




US 20170357478A1

(19) **United States**(12) **Patent Application Publication**
PIERSOL et al.(10) **Pub. No.: US 2017/0357478 A1**(43) **Pub. Date: Dec. 14, 2017**(54) **INTELLIGENT DEVICE ARBITRATION AND CONTROL***G10L 15/22* (2006.01)*G10L 25/84* (2013.01)(71) Applicant: **Apple Inc.**, Cupertino, CA (US)(52) **U.S. Cl.**CPC *G06F 3/167* (2013.01); *G10L 25/84* (2013.01); *G10L 15/30* (2013.01); *G10L 15/22* (2013.01); *G10L 2015/228* (2013.01)(72) Inventors: **Kurt W. PIERSOL**, Cupertino, CA (US); **Ryan M. ORR**, Cupertino, CA (US); **Daniel J. MANDEL**, Tucson, AZ (US)(57) **ABSTRACT**(21) Appl. No.: **15/268,338**(22) Filed: **Sep. 16, 2016****Related U.S. Application Data**

(60) Provisional application No. 62/348,896, filed on Jun. 11, 2016.

Publication Classification(51) **Int. Cl.***G06F 3/16* (2006.01)*G10L 15/30* (2013.01)

This relates to systems and processes for using a virtual assistant to arbitrate among and/or control electronic devices. In one example process, a first electronic device samples an audio input using a microphone. The first electronic device broadcasts a first set of one or more values based on the sampled audio input. Furthermore, the first electronic device receives a second set of one or more values, which are based on the audio input, from a second electronic device. Based on the first set of one or more values and the second set of one or more values, the first electronic device determines whether to respond to the audio input or forego responding to the audio input.






FIG. 1




FIG. 2A

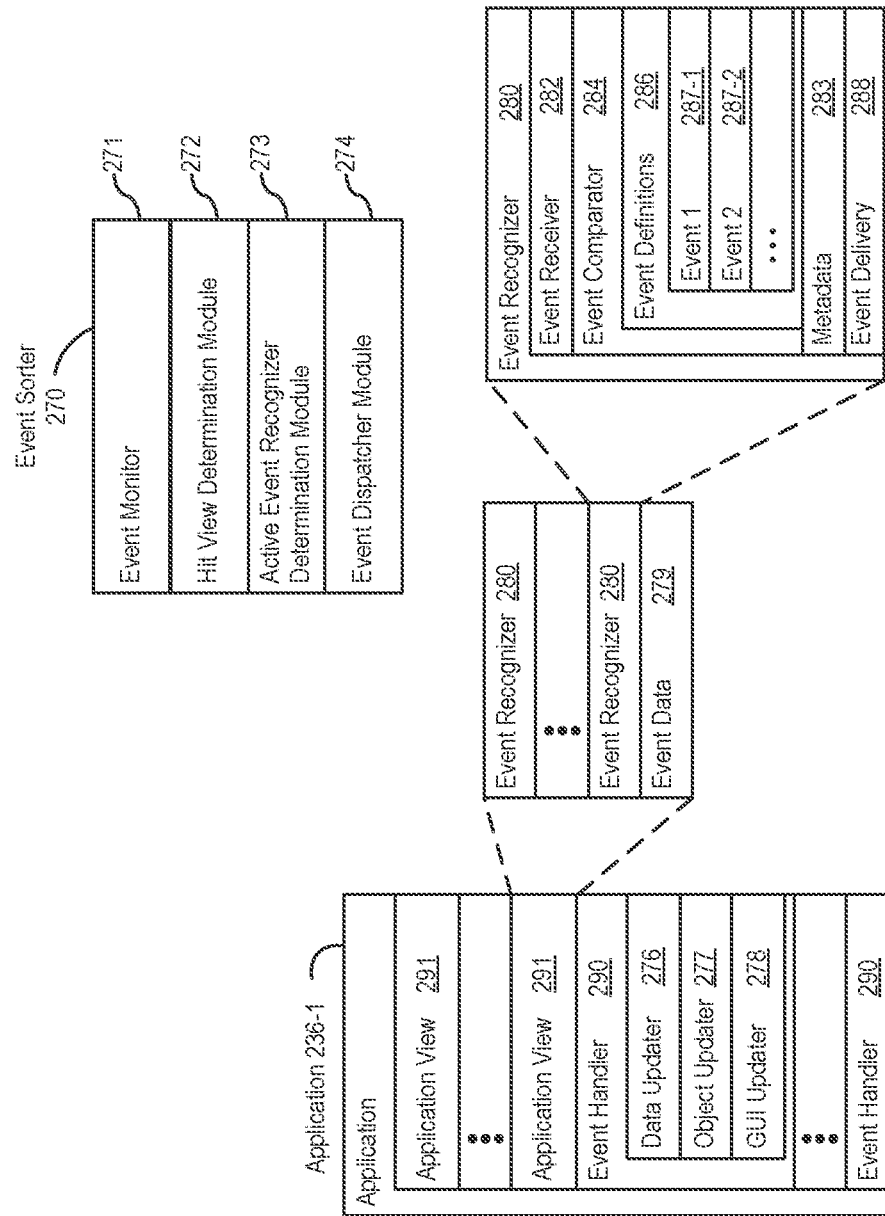


FIG. 2B




FIG. 3

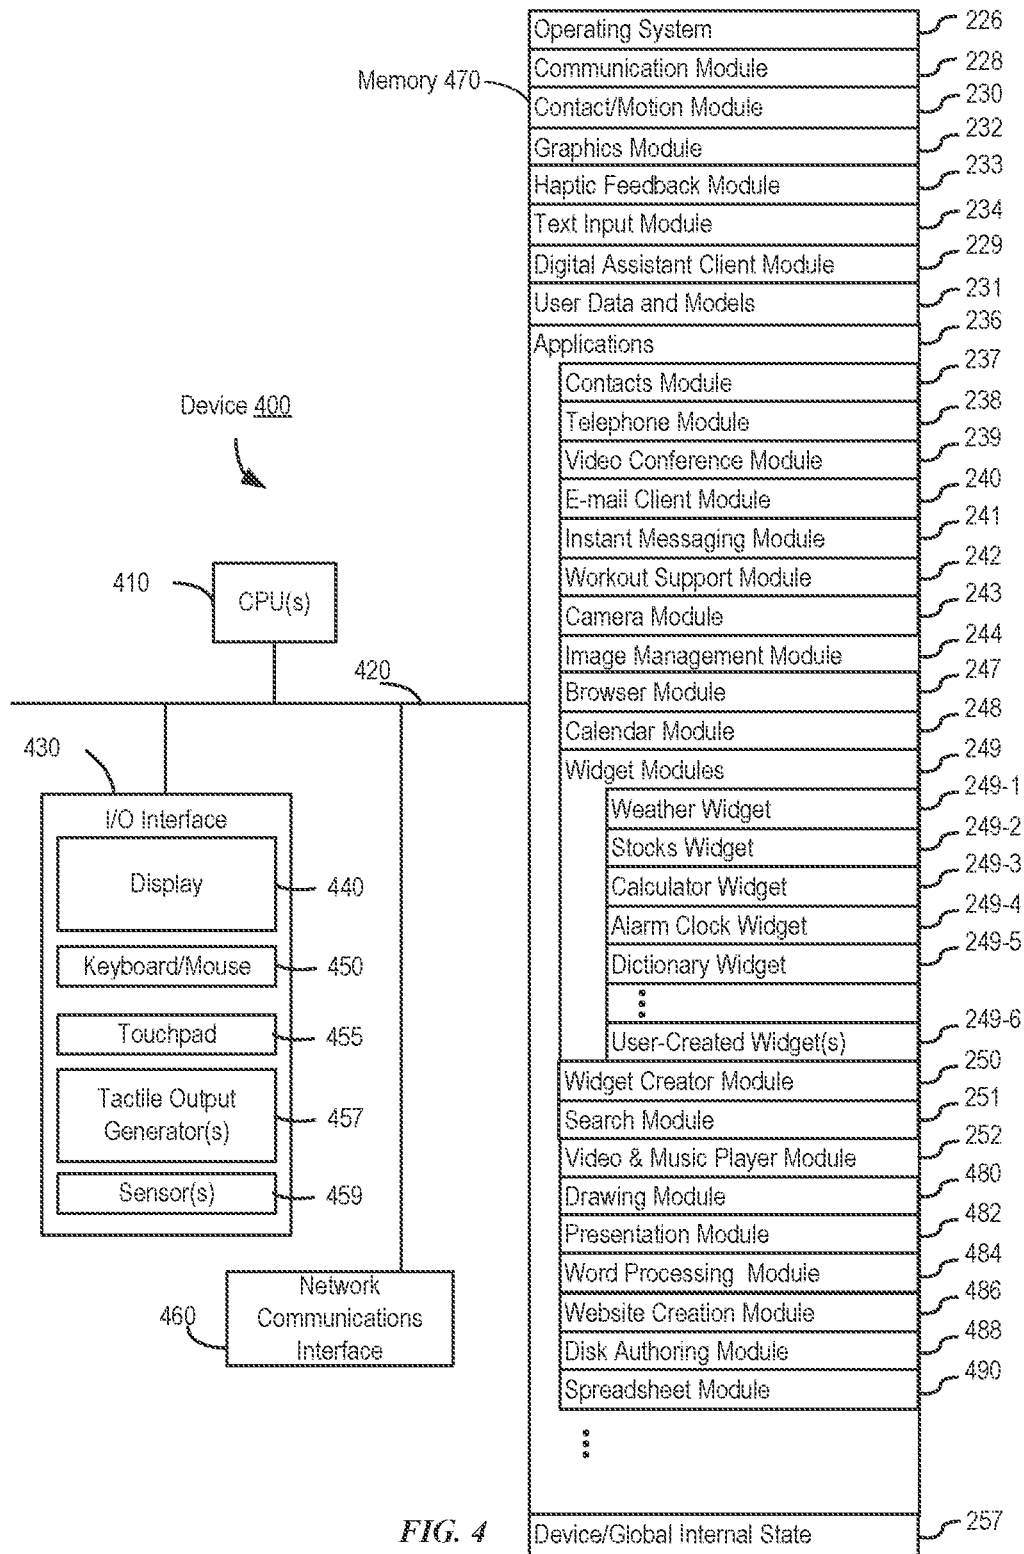


FIG. 4




FIG. 5A





FIG. 5B




FIG. 6A




FIG. 6B

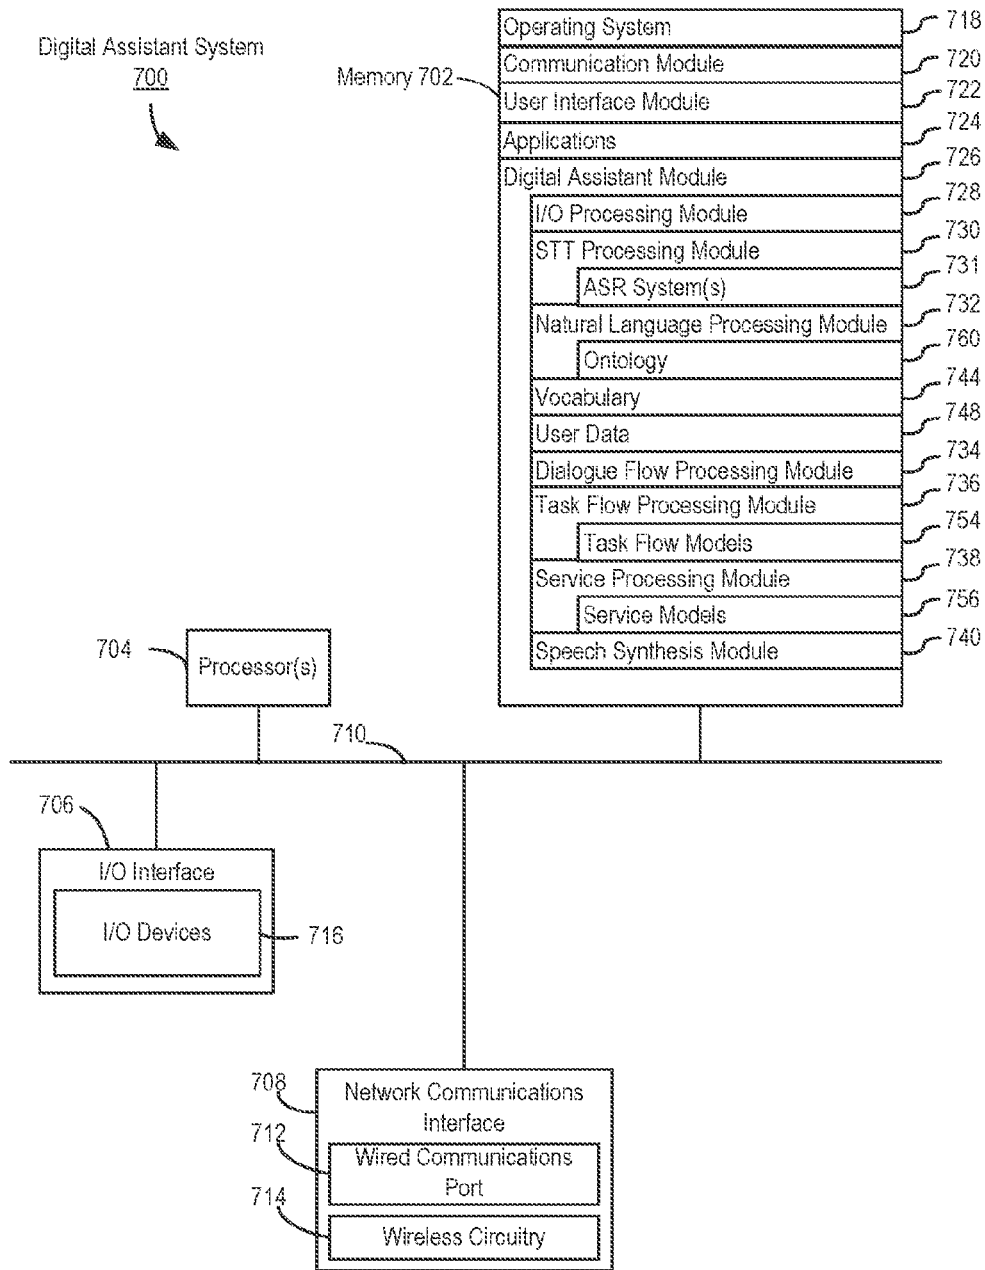


FIG. 7A

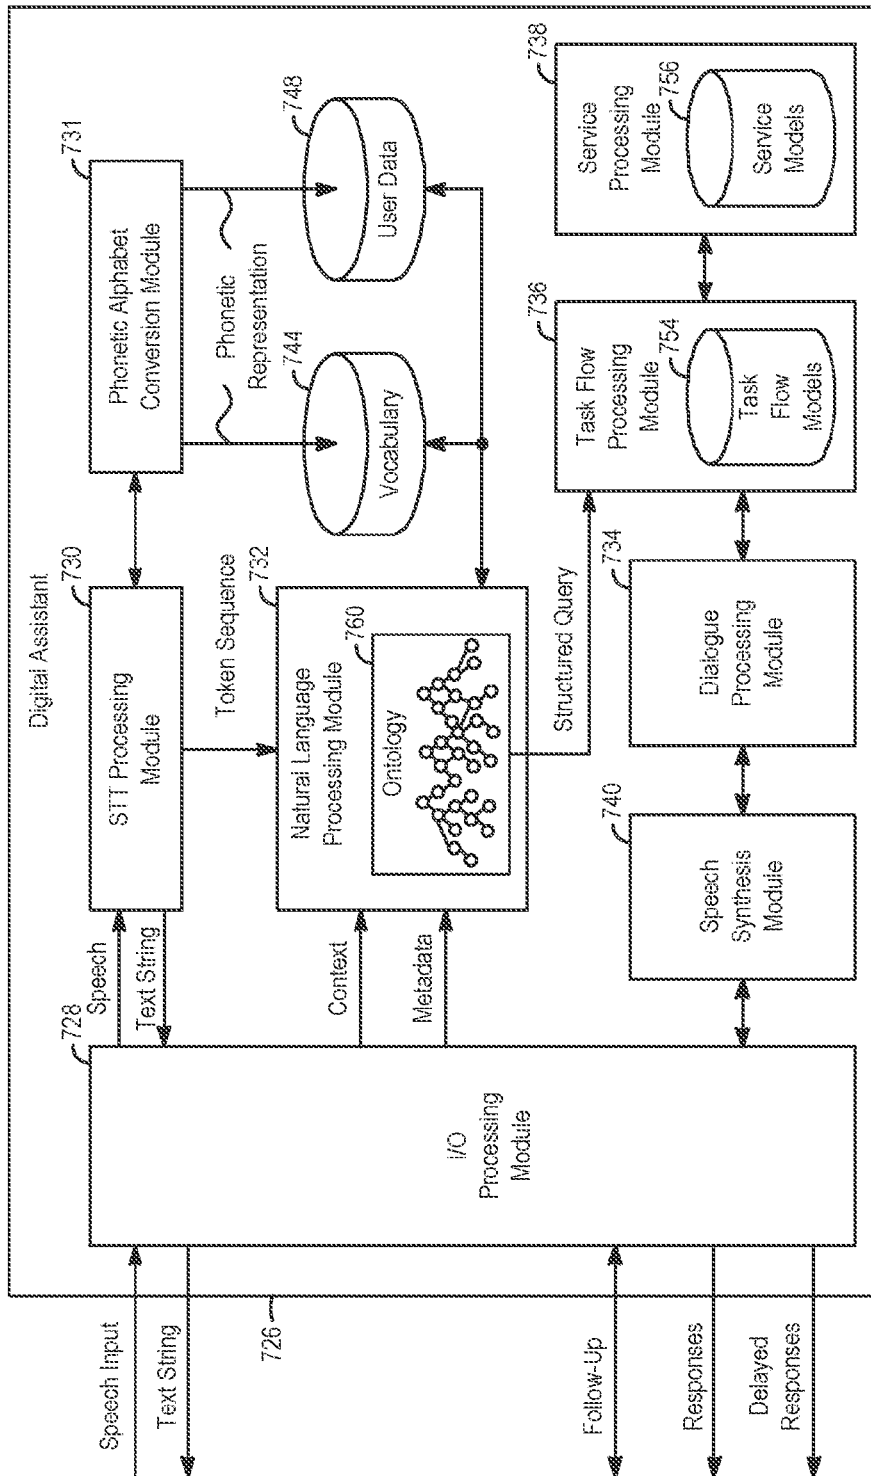


FIG. 7B

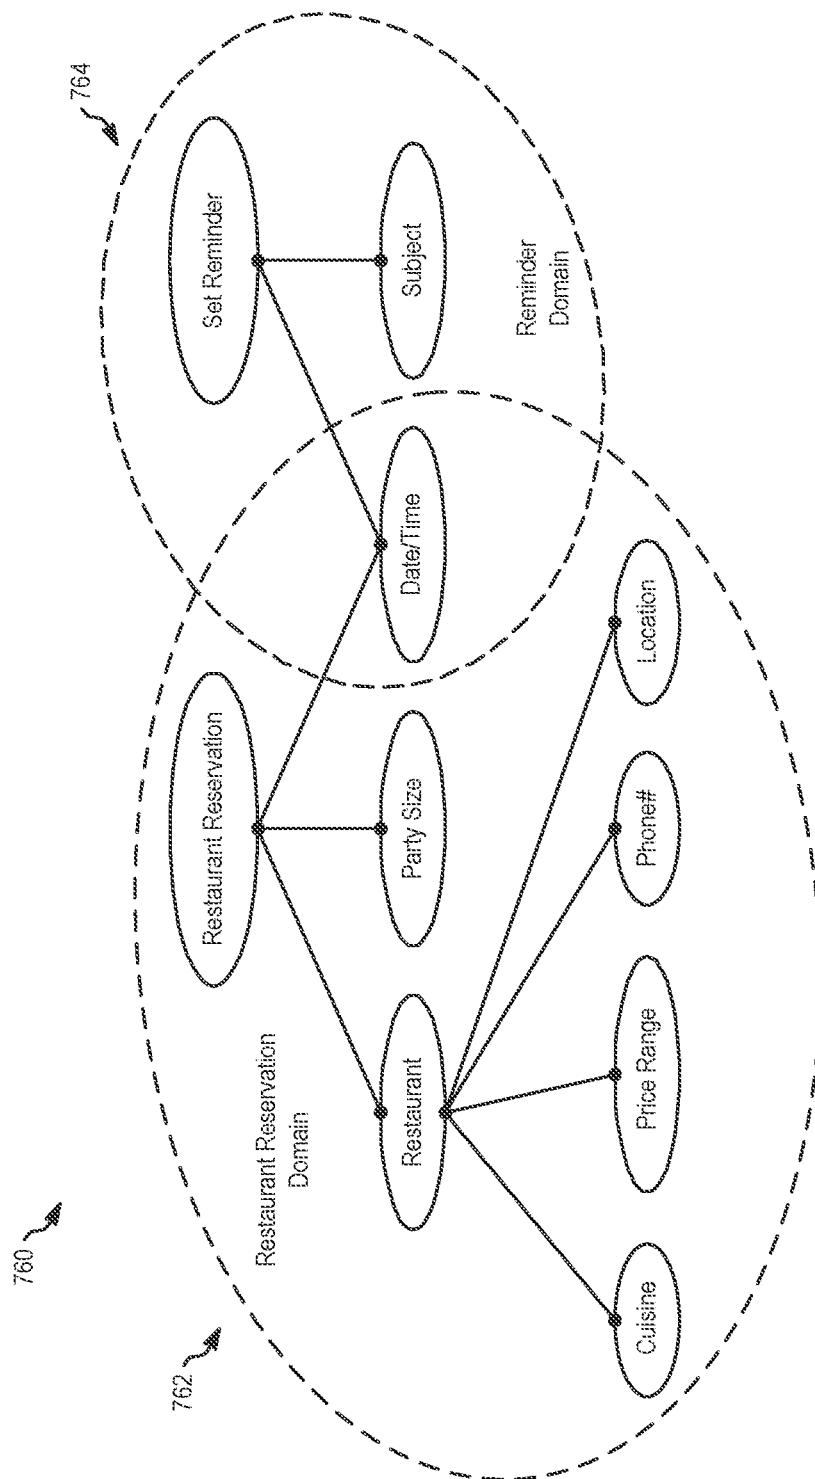


FIG. 7C




FIG. 8A




FIG. 8B




FIG. 8C




FIG. 9A




FIG. 9B




FIG. 9C

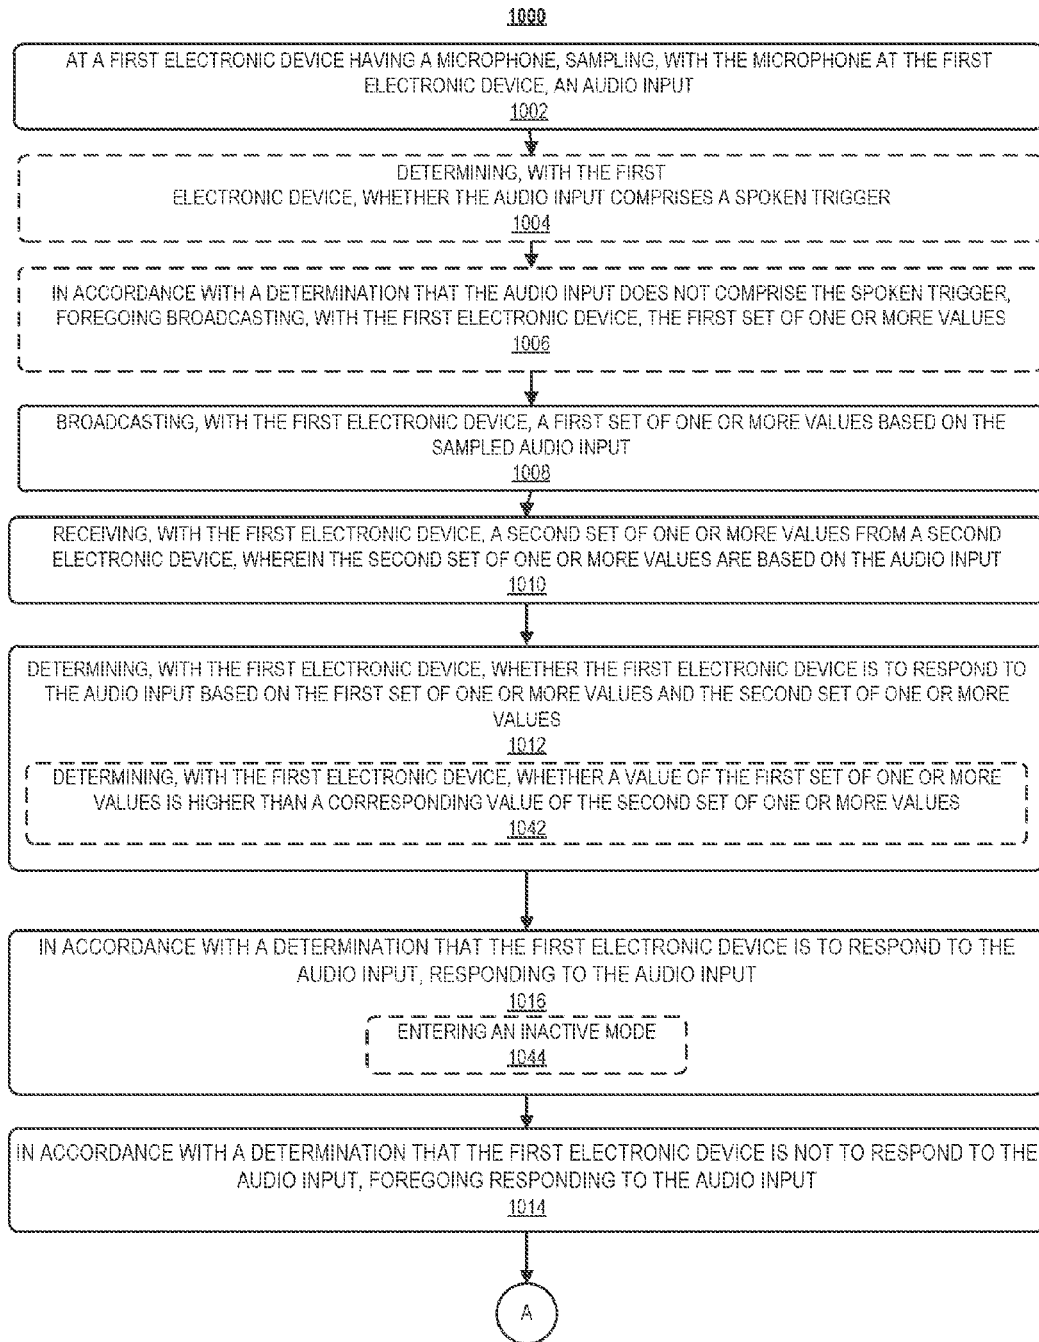


FIG. 10A

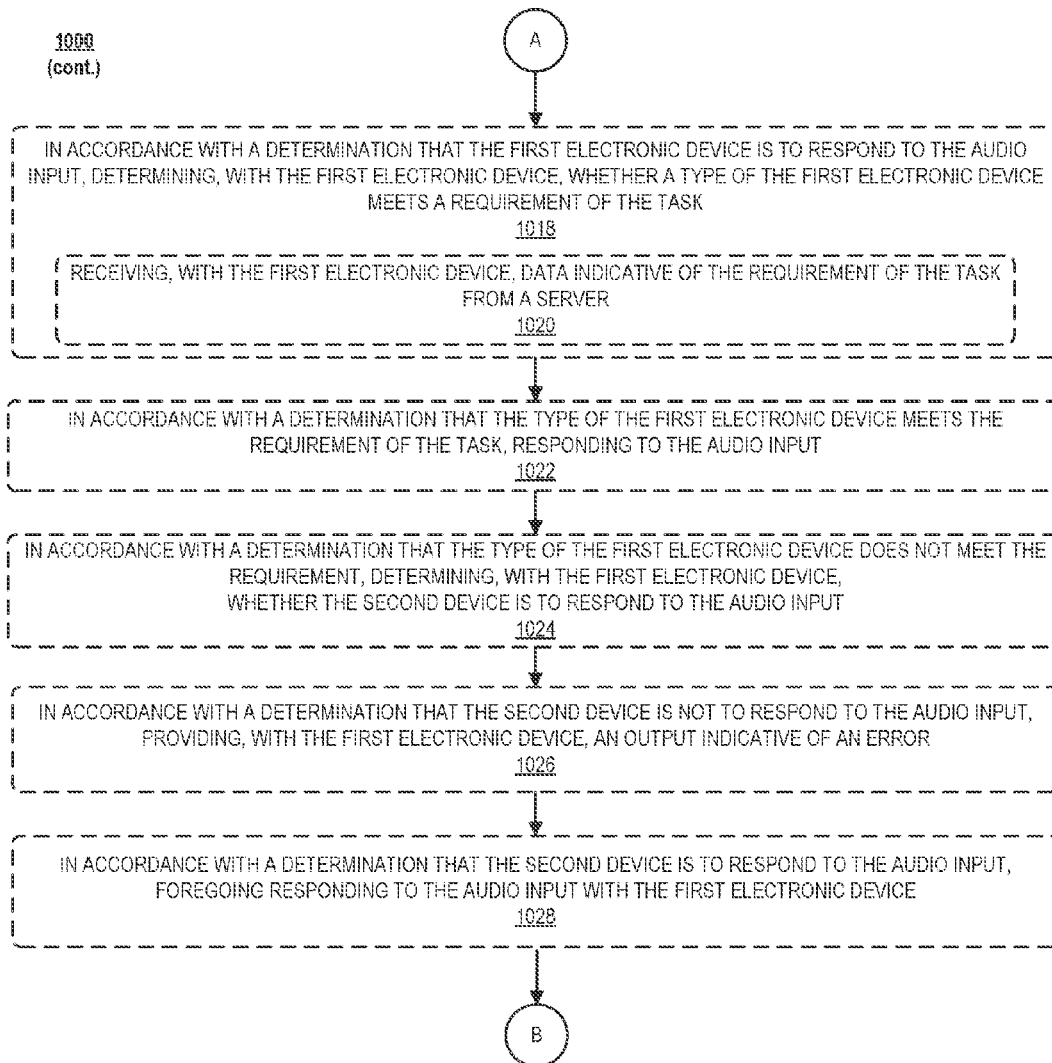


FIG. 10B

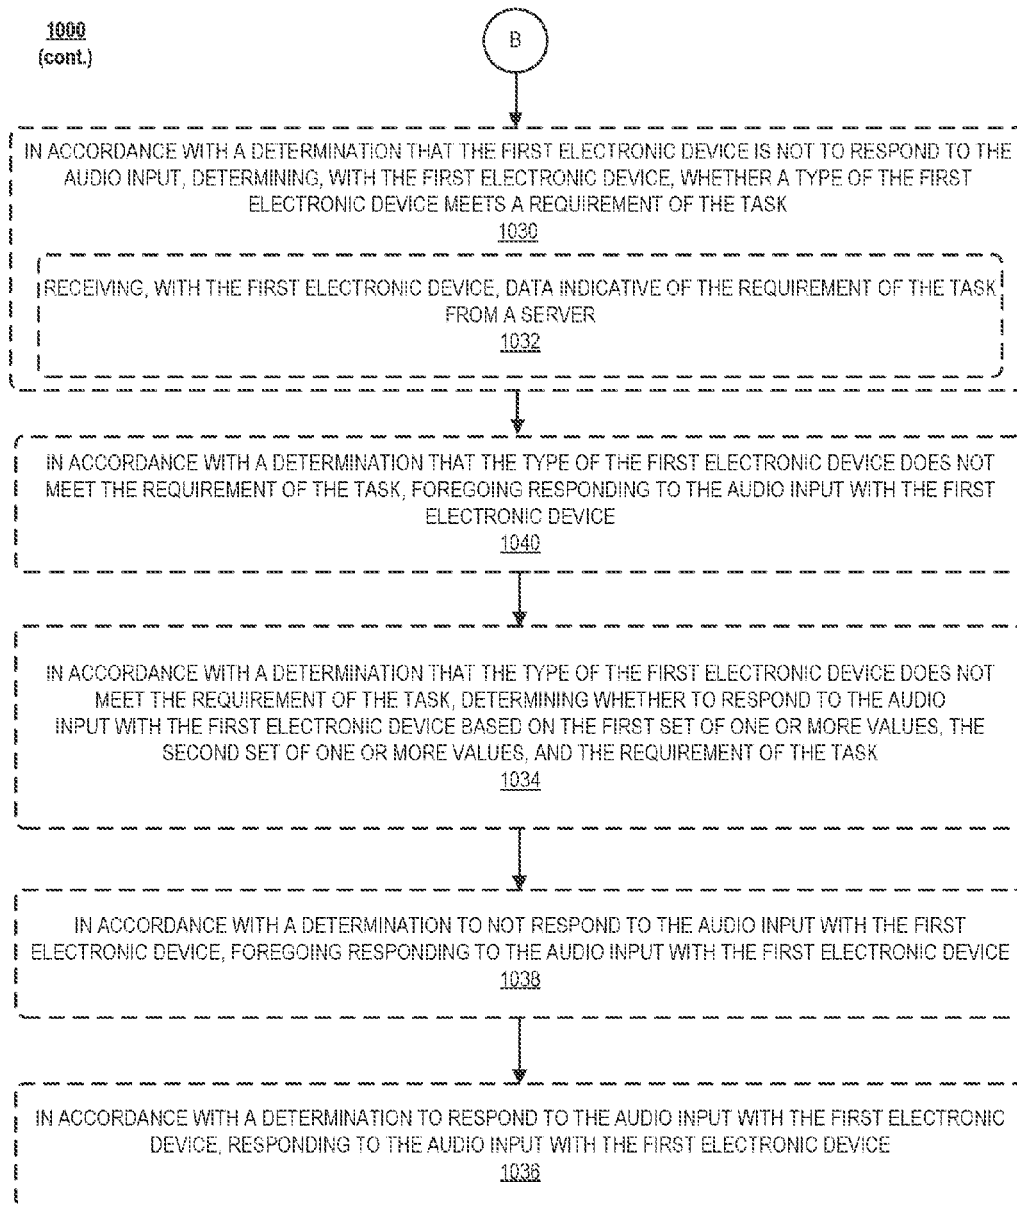


FIG. 10C

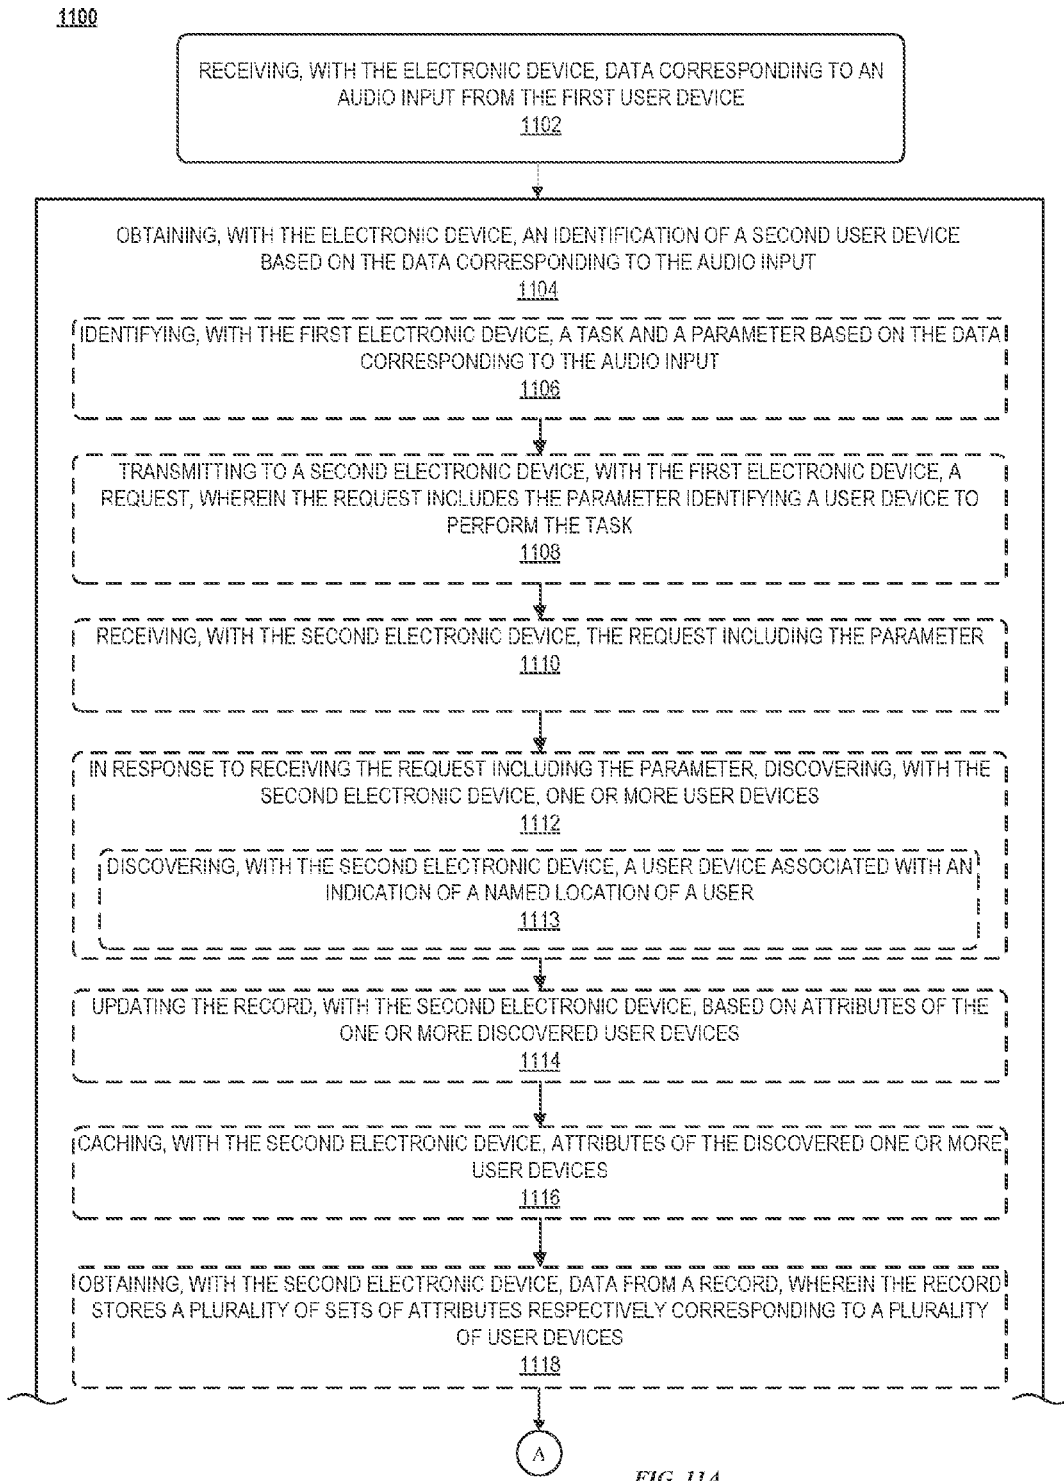


FIG. 11A

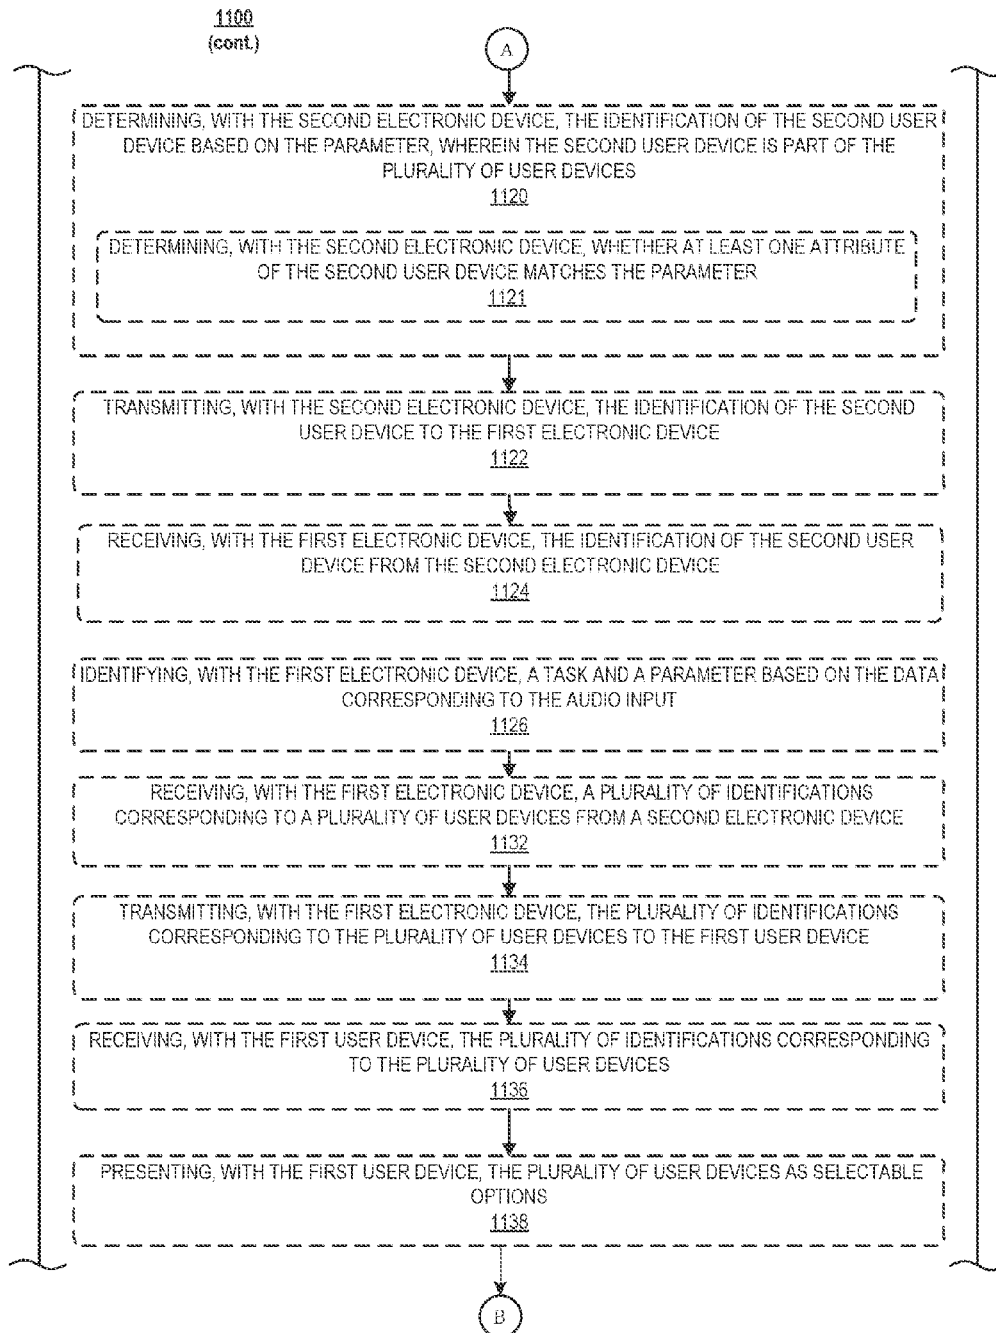
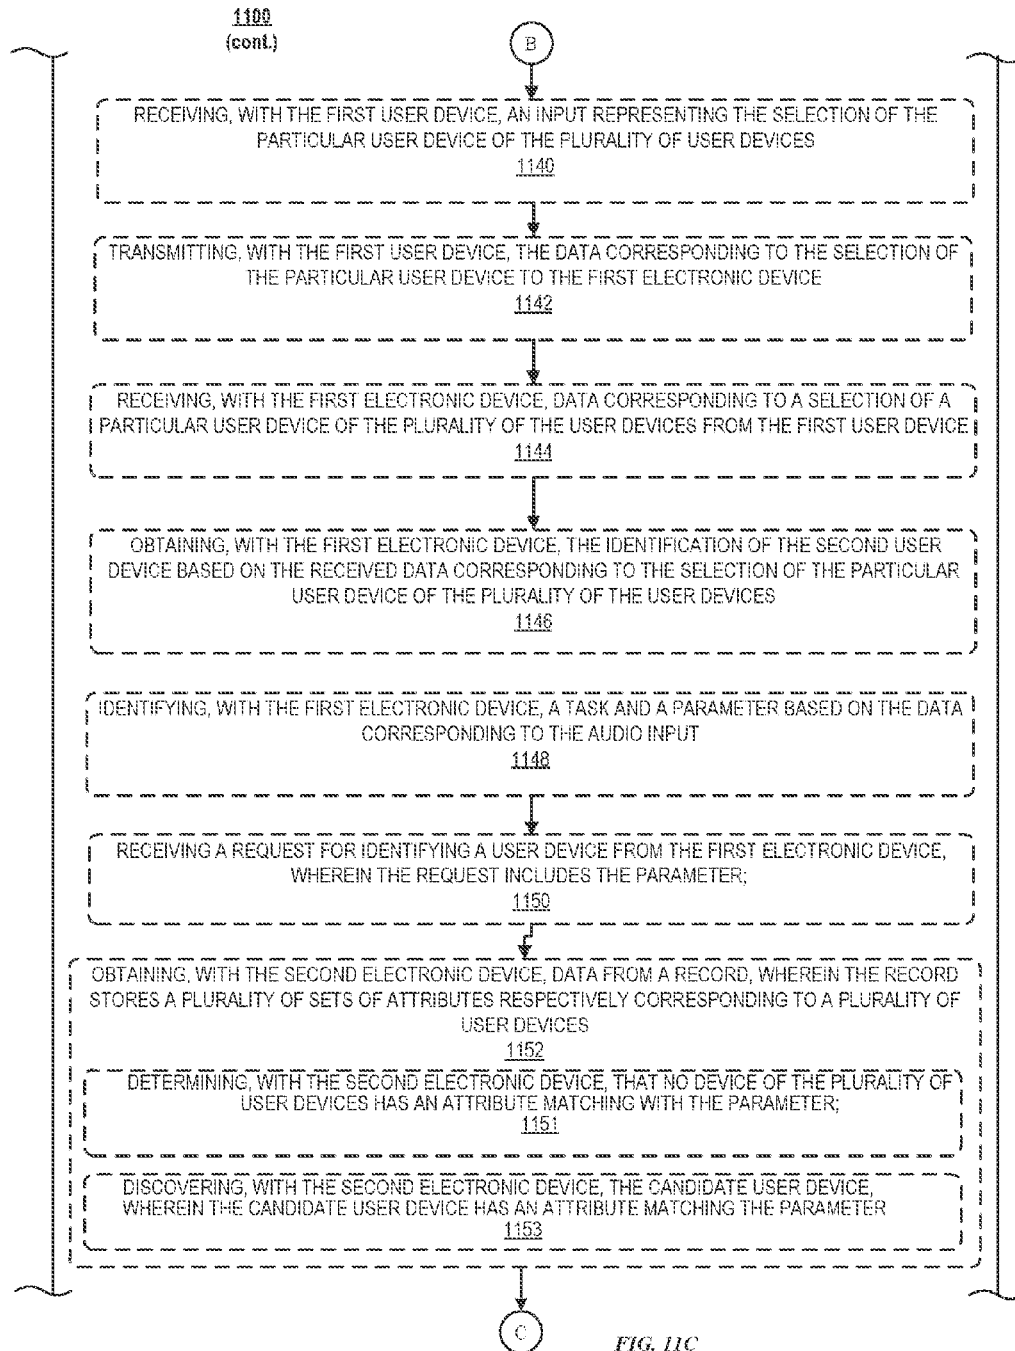


FIG. 11B



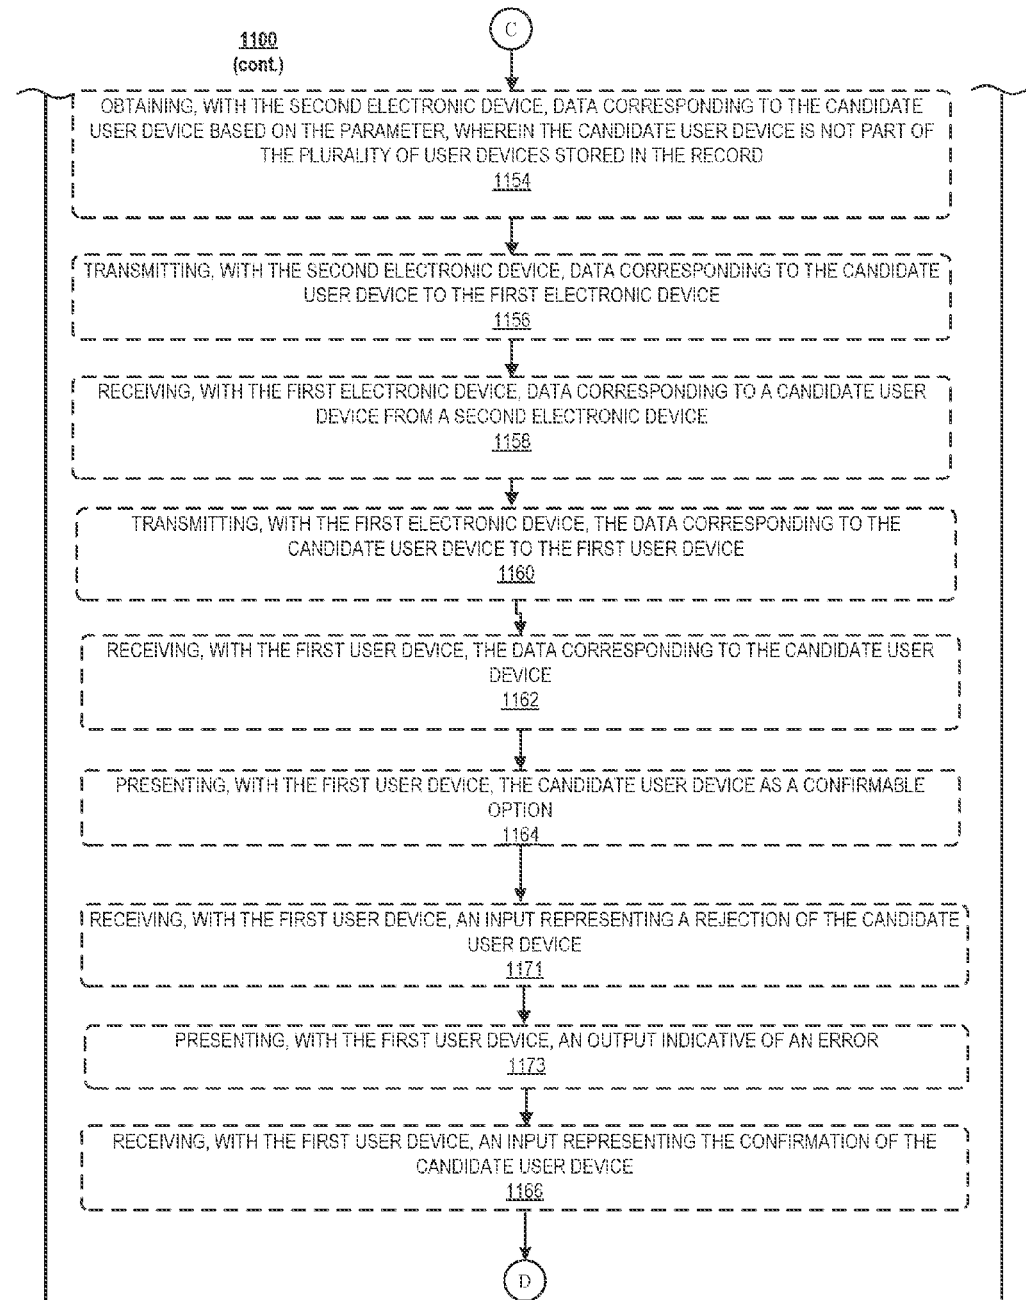


FIG. 11D

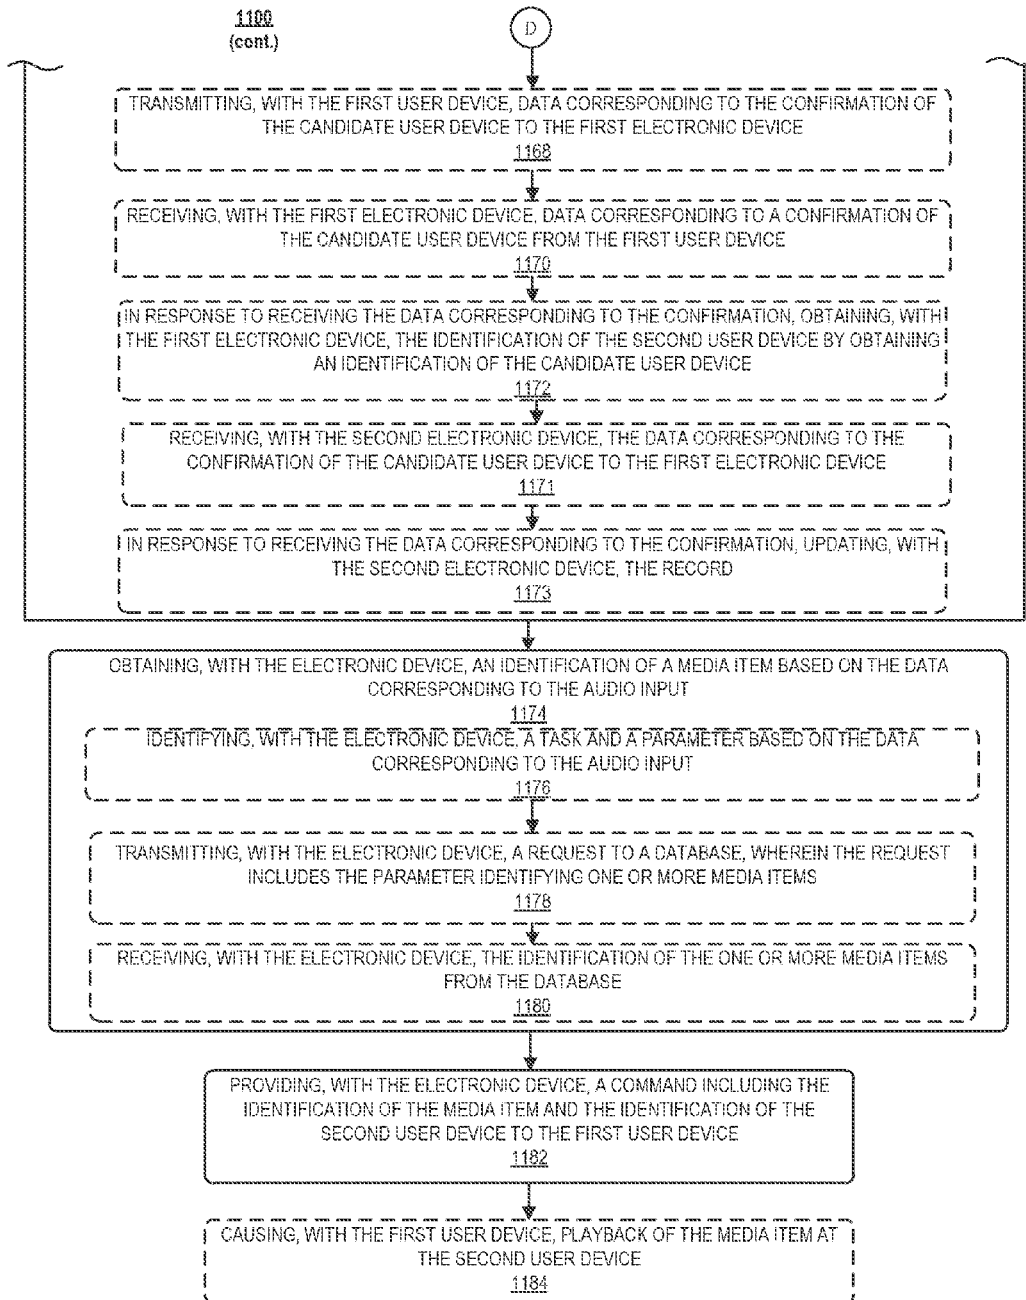


FIG. 11E

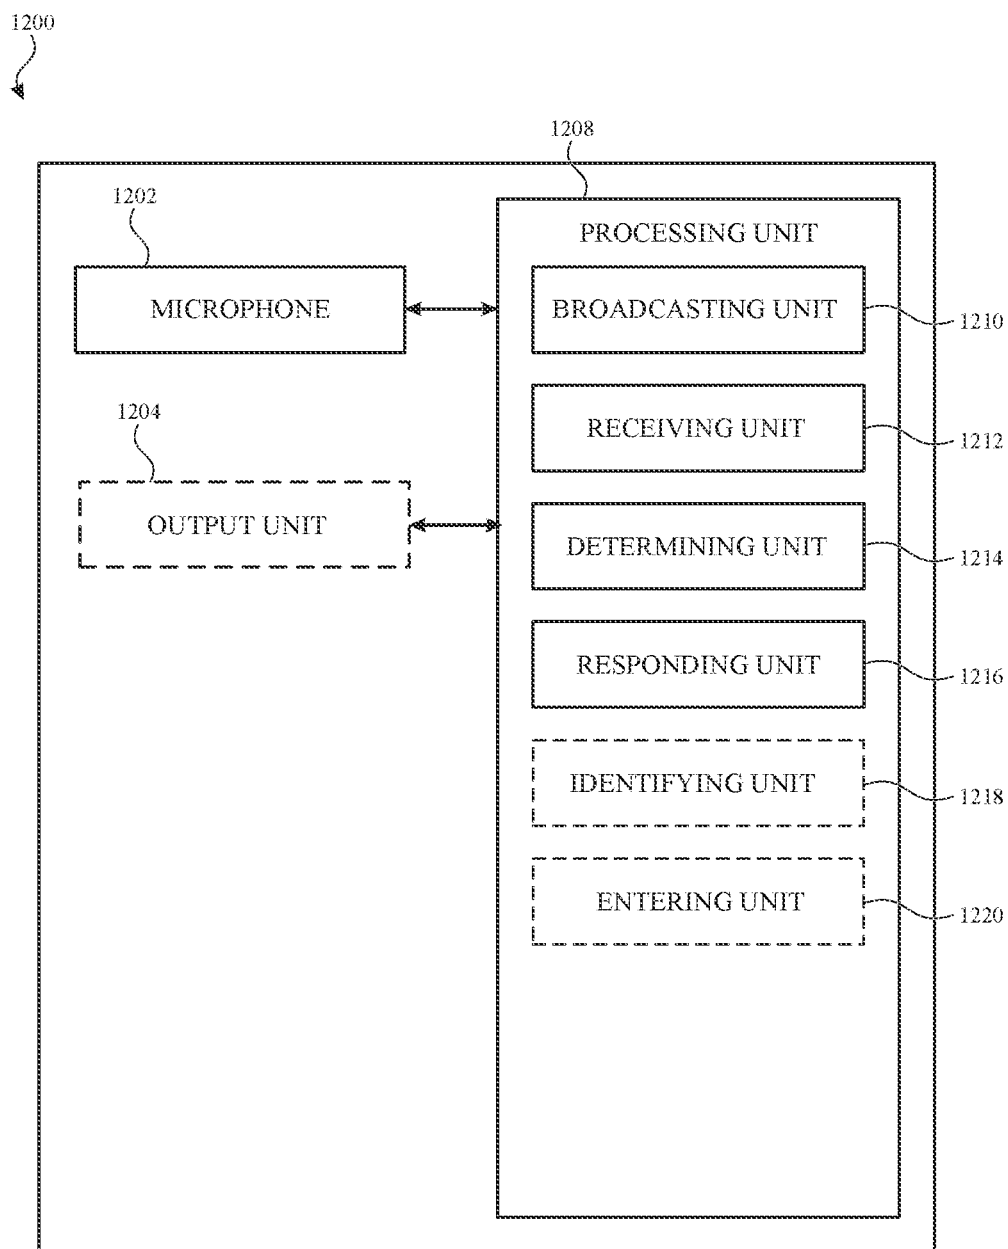


FIG. 12




FIG. 13

INTELLIGENT DEVICE ARBITRATION AND CONTROL

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims priority to U.S. Provisional Patent Application Ser. No. 62/348,896, "INTELLIGENT DEVICE ARBITRATION AND CONTROL," filed on Jun. 11, 2016. The content of this application is hereby incorporated by reference for all purposes.

FIELD

[0002] The present disclosure relates generally to multi-device systems and, more specifically, to intelligent arbitration and control of devices in multi-device systems.

BACKGROUND

[0003] Many contemporary electronic devices offer virtual assistant services to perform various tasks in response to spoken user inputs. In some circumstances, multiple electronic devices having virtual assistant services may concurrently operate in a shared environment. As a result, a user input may cause each of the multiple electronic devices to respond when the user input contains a trigger phrase or command recognized by each of the virtual assistant services on the electronic devices. This in turn may result in a confusing experience for the user, as multiple electronic devices may simultaneously begin to listen and/or prompt for additional input. Further, the multiple electronic devices may perform duplicative or conflicting operations based on the same user input.

[0004] Additionally, a user input can identify one or more particular electronic devices to perform a task. For example, the user may issue a command to a virtual assistant service in order to remotely control home electronic devices by referring to the devices by location ("in my living room"), by a device type ("on my TV"), or by a type of task the device needs to perform ("play the Beatles"). However, the user input, by itself, often does not provide sufficient information for the virtual assistant service to identify, control, and/or manage the electronic devices.

BRIEF SUMMARY

[0005] Example methods are disclosed herein. An example method includes, at a first electronic device having a microphone, sampling, with the microphone at the first electronic device, an audio input; broadcasting, with the first electronic device, a first set of one or more values based on the sampled audio input; receiving, with the first electronic device, a second set of one or more values from a second electronic device, where the second set of one or more values are based on the audio input; and determining, with the first electronic device, whether the first electronic device is to respond to the audio input based on the first set of one or more values and the second set of one or more values. In accordance with a determination that the first electronic device is to respond to the audio input, the first electronic device responds to the audio input. In accordance with a determination that the first electronic device is not to respond to the audio input, the first electronic device foregoes responding to the audio input.

[0006] Example non-transitory computer readable storage media are disclosed herein. An example non-transitory com-

puter readable medium stores one or more programs. The one or more programs comprise instructions, which when executed by one or more processors of an electronic device with a microphone, cause the electronic device to sample, with the microphone at the electronic device, an audio input; broadcast a first set of one or more values based on the sampled audio input; receive a second set of one or more values from a second electronic device, where the second set of one or more values are based on the audio input; and determine whether the electronic device is to respond to the audio input based on the first set of one or more values and the second set of one or more values. The instructions further cause the electronic device to, in accordance with a determination that the electronic device is to respond to the audio input, respond to the audio input; in accordance with a determination that the electronic device is not to respond to the audio input, forego responding to the audio input.

[0007] Example devices are disclosed herein. An example device comprises a microphone; one or more processors; a memory; and one or more programs, where the one or more programs are stored in the memory and configured to be executed by the one or more processors. The one or more programs include instructions for sampling, with the microphone at the electronic device, an audio input; broadcasting a first set of one or more values based on the sampled audio input; receiving a second set of one or more values from a second electronic device, where the second set of one or more values are based on the audio input; and determining whether the electronic device is to respond to the audio input based on the first set of one or more values and the second set of one or more values. The one or more programs further include instructions for, in accordance with a determination that the electronic device is to respond to the audio input, responding to the audio input; and in accordance with a determination that the electronic device is not to respond to the audio input, foregoing responding to the audio input.

[0008] An example electronic device comprises a microphone; means for sampling, with the microphone at the electronic device, an audio input; means for broadcasting a first set of one or more values based on the sampled audio input; means for receiving a second set of one or more values from a second electronic device, where the second set of one or more values are based on the audio input; means for determining whether the electronic device is to respond to the audio input based on the first set of one or more values and the second set of one or more values; means for, in accordance with a determination that the electronic device is to respond to the audio input, responding to the audio input; and means for, in accordance with a determination that the electronic device is not to respond to the audio input, foregoing responding to the audio input.

[0009] An example method for controlling user devices using a virtual assistant on a first user device comprises, at an electronic device, receiving, with the electronic device, data corresponding to an audio input from the first user device; obtaining, with the electronic device, an identification of a second user device based on the data corresponding to the audio input; obtaining, with the electronic device, an identification of a media item based on the data corresponding to the audio input; and providing, with the electronic device, a command including the identification of the media item and the identification of the second user device to the first user device.

[0010] An example non-transitory computer readable storage medium stores one or more programs. The one or more programs comprise instructions, which when executed by one or more processors of an electronic device, cause the electronic device to receive data corresponding to an audio input from a first user device; obtain an identification of a second user device based on the data corresponding to the audio input; obtain an identification of a media item based on the data corresponding to the audio input; and provide a command including the identification of the media item and the identification of the second user device to the first user device.

[0011] An example electronic device comprises one or more processors; a memory; and one or more programs, where the one or more programs are stored in the memory and configured to be executed by the one or more processors, the one or more programs including instructions for receiving data corresponding to an audio input from a first user device; obtaining an identification of a second user device based on the data corresponding to the audio input; obtaining an identification of a media item based on the data corresponding to the audio input; and providing a command including the identification of the media item and the identification of the second user device to the first user device.

[0012] An example electronic device comprises means for receiving data corresponding to an audio input from a first user device; means for obtaining an identification of a second user device based on the data corresponding to the audio input; means for obtaining an identification of a media item based on the data corresponding to the audio input; and means for providing a command including the identification of the media item and the identification of the second user device to the first user device.

DESCRIPTION OF THE FIGURES

[0013] For a better understanding of the various described embodiments, reference should be made to the Detailed Description below, in conjunction with the following drawings in which like reference numerals refer to corresponding parts throughout the figures.

[0014] FIG. 1 is a block diagram illustrating a system and environment for implementing a digital assistant according to various examples.

[0015] FIG. 2A is a block diagram illustrating a portable multifunction device implementing the client-side portion of a digital assistant in accordance with some embodiments.

[0016] FIG. 2B is a block diagram illustrating exemplary components for event handling according to various examples.

[0017] FIG. 3 illustrates a portable multifunction device implementing the client-side portion of a digital assistant according to various examples.

[0018] FIG. 4 is a block diagram of an exemplary multifunction device with a display and a touch-sensitive surface according to various examples.

[0019] FIG. 5A illustrates an exemplary user interface for a menu of applications on a portable multifunction device according to various examples.

[0020] FIG. 5B illustrates an exemplary user interface for a multifunction device with a touch-sensitive surface that is separate from the display according to various examples.

[0021] FIG. 6A illustrates a personal electronic device according to various examples.

[0022] FIG. 6B is a block diagram illustrating a personal electronic device according to various examples.

[0023] FIG. 7A is a block diagram illustrating a digital assistant system or a server portion thereof according to various examples.

[0024] FIG. 7B illustrates the functions of the digital assistant shown in FIG. 7A according to various examples.

[0025] FIG. 7C illustrates a portion of an ontology according to various examples.

[0026] FIGS. 8A-C illustrate a plurality of electronic devices according to various examples.

[0027] FIGS. 9A-9C illustrate exemplary system and environment for controlling electronic devices according to various examples.

[0028] FIGS. 10A-C illustrate an exemplary process for device arbitration according to various examples.

[0029] FIGS. 11A-E illustrate an exemplary process for device control according to various examples.

[0030] FIG. 12 illustrates a functional block diagram of an electronic device according to various examples.

[0031] FIG. 13 illustrates functional block diagrams of electronic devices of a system according to various examples.

DETAILED DESCRIPTION

[0032] In the following description of the disclosure and embodiments, reference is made to the accompanying drawings in which it is shown by way of illustration of specific embodiments that can be practiced. It is to be understood that other embodiments and examples can be practiced and changes can be made without departing from the scope of the disclosure.

[0033] Techniques for intelligent device arbitration and control are desirable. As described herein, described techniques improve the user's ability to interact with multiple electronic devices, thereby enhancing productivity. Further, such techniques can reduce computational demand and battery power otherwise consumed as a result of redundant responses by electronic devices to user inputs.

[0034] Although the following description uses terms "first," "second," etc. to describe various elements, these elements should not be limited by the terms. These terms are only used to distinguish one element from another. For example, a first input could be termed a second input, and, similarly, a second input could be termed a first input, without departing from the scope of the various described examples. The first input and the second input can both be outputs and, in some cases, can be separate and different inputs.

[0035] The terminology used in the description of the various described examples herein is for the purpose of describing particular examples only and is not intended to be limiting. As used in the description of the various described examples and the appended claims, the singular forms "a," "an," and "the" are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will also be understood that the term "and/or" as used herein refers to and encompasses any and all possible combinations of one or more of the associated listed items. It will be further understood that the terms "includes," "including," "comprises," and/or "comprising," when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not

preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof.

[0036] The term “if” may be construed to mean “when” or “upon” or “in response to determining” or “in response to detecting,” depending on the context. Similarly, the phrase “if it is determined” or “if [a stated condition or event] is detected” may be construed to mean “upon determining” or “in response to determining” or “upon detecting [the stated condition or event]” or “in response to detecting [the stated condition or event],” depending on the context.

1. System and Environment

[0037] FIG. 1 illustrates a block diagram of system 100 according to various examples. In some examples, system 100 can implement a digital assistant. The terms “digital assistant,” “virtual assistant,” “intelligent automated assistant,” or “automatic digital assistant” can refer to any information processing system that interprets natural language input in spoken and/or textual form to infer user intent, and performs actions based on the inferred user intent. For example, to act on an inferred user intent, the system can perform one or more of the following: identifying a task flow with steps and parameters designed to accomplish the inferred user intent, inputting specific requirements from the inferred user intent into the task flow; executing the task flow by invoking programs, methods, services, APIs, or the like; and generating output responses to the user in an audible (e.g., speech) and/or visual form.

[0038] Specifically, a digital assistant can be capable of accepting a user request at least partially in the form of a natural language command, request, statement, narrative, and/or inquiry. Typically, the user request can seek either an informational answer or performance of a task by the digital assistant. A satisfactory response to the user request can be a provision of the requested informational answer, a performance of the requested task, or a combination of the two. For example, a user can ask the digital assistant a question, such as “Where am I right now?” Based on the user’s current location, the digital assistant can answer, “You are in Central Park near the west gate.” The user can also request the performance of a task, for example, “Please invite my friends to my girlfriend’s birthday party next week.” In response, the digital assistant can acknowledge the request by saying “Yes, right away,” and then send a suitable calendar invite on behalf of the user to each of the user’s friends listed in the user’s electronic address book. During performance of a requested task, the digital assistant can sometimes interact with the user in a continuous dialogue involving multiple exchanges of information over an extended period of time. There are numerous other ways of interacting with a digital assistant to request information or performance of various tasks. In addition to providing verbal responses and taking programmed actions, the digital assistant can also provide responses in other visual or audio forms, e.g., as text, alerts, music, videos, animations, etc.

[0039] As shown in FIG. 1, in some examples, a digital assistant can be implemented according to a client-server model. The digital assistant can include client-side portion 102 (hereafter “DA client 102”) executed on user device 104 and server-side portion 106 (hereafter “DA server 106”) executed on server system 108. DA client 102 can communicate with DA server 106 through one or more networks 110. DA client 102 can provide client-side functionalities

such as user-facing input and output processing and communication with DA server 106. DA server 106 can provide server-side functionalities for any number of DA clients 102 each residing on a respective user device 104.

[0040] In some examples, DA server 106 can include client-facing I/O interface 112, one or more processing modules 114, data and models 116, and I/O interface to external services 118. The client-facing I/O interface 112 can facilitate the client-facing input and output processing for DA server 106. One or more processing modules 114 can utilize data and models 116 to process speech input and determine the user’s intent based on natural language input. Further, one or more processing modules 114 perform task execution based on inferred user intent. In some examples, DA server 106 can communicate with external services 120 through network(s) 110 for task completion or information acquisition. I/O interface to external services 118 can facilitate such communications.

[0041] User device 104 can be any suitable electronic device. For example, user devices can be a portable multifunctional device (e.g., device 200, described below with reference to FIG. 2A), a multifunctional device (e.g., device 400, described below with reference to FIG. 4), or a personal electronic device (e.g., device 600, described below with reference to FIG. 6A-B.) A portable multifunctional device can be, for example, a mobile telephone that also contains other functions, such as PDA and/or music player functions. Specific examples of portable multifunction devices can include the iPhone®, iPod Touch®, and iPad® devices from Apple Inc. of Cupertino, Calif. Other examples of portable multifunction devices can include, without limitation, laptop or tablet computers. Further, in some examples, user device 104 can be a non-portable multifunctional device. In particular, user device 104 can be a desktop computer, a game console, a television, or a television set-top box. In some examples, user device 104 can include a touch-sensitive surface (e.g., touch screen displays and/or touchpads). Further, user device 104 can optionally include one or more other physical user-interface devices, such as a physical keyboard, a mouse, and/or a joystick. Various examples of electronic devices, such as multifunctional devices, are described below in greater detail.

[0042] Examples of communication network(s) 110 can include local area networks (LAN) and wide area networks (WAN), e.g., the Internet. Communication network(s) 110 can be implemented using any known network protocol, including various wired or wireless protocols, such as, for example, Ethernet, Universal Serial Bus (USB), FIREWIRE, Global System for Mobile Communications (GSM), Enhanced Data GSM Environment (EDGE), code division multiple access (CDMA), time division multiple access (TDMA), Bluetooth, Wi-Fi, voice over Internet Protocol (VoIP), Wi-MAX, or any other suitable communication protocol.

[0043] Server system 108 can be implemented on one or more standalone data processing apparatus or a distributed network of computers. In some examples, server system 108 can also employ various virtual devices and/or services of third-party service providers (e.g., third-party cloud service providers) to provide the underlying computing resources and/or infrastructure resources of server system 108.

[0044] In some examples, user device 104 can communicate with DA server 106 via second user device 122. Second user device 122 can be similar or identical to user device

104. For example, second user device **122** can be similar to devices **200**, **400**, or **600** described below with reference to FIGS. **2A**, **4**, and **6A-B**. User device **104** can be configured to communicatively couple to second user device **122** via a direct communication connection, such as Bluetooth, NFC, BTLE, or the like, or via a wired or wireless network, such as a local Wi-Fi network. In some examples, second user device **122** can be configured to act as a proxy between user device **104** and DA server **106**. For example, DA client **102** of user device **104** can be configured to transmit information (e.g., a user request received at user device **104**) to DA server **106** via second user device **122**. DA server **106** can process the information and return relevant data (e.g., data content responsive to the user request) to user device **104** via second user device **122**.

[0045] In some examples, user device **104** can be configured to communicate abbreviated requests for data to second user device **122** to reduce the amount of information transmitted from user device **104**. Second user device **122** can be configured to determine supplemental information to add to the abbreviated request to generate a complete request to transmit to DA server **106**. This system architecture can advantageously allow user device **104** having limited communication capabilities and/or limited battery power (e.g., a watch or a similar compact electronic device) to access services provided by DA server **106** by using second user device **122**, having greater communication capabilities and/or battery power (e.g., a mobile phone, laptop computer, tablet computer, or the like), as a proxy to DA server **106**. While only two user devices **104** and **122** are shown in FIG. **1**, it will be appreciated that system **100** can include any number and type of user devices configured in this proxy configuration to communicate with DA server system **106**.

[0046] Although the digital assistant shown in FIG. **1** can include both a client-side portion (e.g., DA client **102**) and a server-side portion (e.g., DA server **106**), in some examples, the functions of a digital assistant can be implemented as a standalone application installed on a user device. In addition, the divisions of functionalities between the client and server portions of the digital assistant can vary in different implementations. For instance, in some examples, the DA client can be a thin-client that provides only user-facing input and output processing functions, and delegates all other functionalities of the digital assistant to a backend server.

2. Electronic Devices

[0047] Attention is now directed toward embodiments of electronic devices for implementing the client-side portion of a digital assistant. FIG. **2A** is a block diagram illustrating portable multifunction device **200** with touch-sensitive display system **212** in accordance with some embodiments. Touch-sensitive display **212** is sometimes called a “touch screen” for convenience and is sometimes known as or called a “touch-sensitive display system.” Device **200** includes memory **202** (which optionally includes one or more computer-readable storage mediums), memory controller **222**, one or more processing units (CPUs) **220**, peripherals interface **218**, RF circuitry **208**, audio circuitry **210**, speaker **211**, microphone **213**, input/output (I/O) subsystem **206**, other input control devices **216**, and external port **224**. Device **200** optionally includes one or more optical sensors **264**. Device **200** optionally includes one or more contact intensity sensors **265** for detecting intensity of

contacts on device **200** (e.g., a touch-sensitive surface such as touch-sensitive display system **212** of device **200**). Device **200** optionally includes one or more tactile output generators **267** for generating tactile outputs on device **200** (e.g., generating tactile outputs on a touch-sensitive surface such as touch-sensitive display system **212** of device **200** or touchpad **455** of device **400**). These components optionally communicate over one or more communication buses or signal lines **203**.

[0048] As used in the specification and claims, the term “intensity” of a contact on a touch-sensitive surface refers to the force or pressure (force per unit area) of a contact (e.g., a finger contact) on the touch-sensitive surface, or to a substitute (proxy) for the force or pressure of a contact on the touch-sensitive surface. The intensity of a contact has a range of values that includes at least four distinct values and more typically includes hundreds of distinct values (e.g., at least **256**). Intensity of a contact is, optionally, determined (or measured) using various approaches and various sensors or combinations of sensors. For example, one or more force sensors underneath or adjacent to the touch-sensitive surface are, optionally, used to measure force at various points on the touch-sensitive surface. In some implementations, force measurements from multiple force sensors are combined (e.g., a weighted average) to determine an estimated force of a contact. Similarly, a pressure-sensitive tip of a stylus is, optionally, used to determine a pressure of the stylus on the touch-sensitive surface. Alternatively, the size of the contact area detected on the touch-sensitive surface and/or changes thereto, the capacitance of the touch-sensitive surface proximate to the contact and/or changes thereto, and/or the resistance of the touch-sensitive surface proximate to the contact and/or changes thereto are, optionally, used as a substitute for the force or pressure of the contact on the touch-sensitive surface. In some implementations, the substitute measurements for contact force or pressure are used directly to determine whether an intensity threshold has been exceeded (e.g., the intensity threshold is described in units corresponding to the substitute measurements). In some implementations, the substitute measurements for contact force or pressure are converted to an estimated force or pressure, and the estimated force or pressure is used to determine whether an intensity threshold has been exceeded (e.g., the intensity threshold is a pressure threshold measured in units of pressure). Using the intensity of a contact as an attribute of a user input allows for user access to additional device functionality that may otherwise not be accessible by the user on a reduced-size device with limited real estate for displaying affordances (e.g., on a touch-sensitive display) and/or receiving user input (e.g., via a touch-sensitive display, a touch-sensitive surface, or a physical/mechanical control such as a knob or a button).

[0049] As used in the specification and claims, the term “tactile output” refers to physical displacement of a device relative to a previous position of the device, physical displacement of a component (e.g., a touch-sensitive surface) of a device relative to another component (e.g., housing) of the device, or displacement of the component relative to a center of mass of the device that will be detected by a user with the user’s sense of touch. For example, in situations where the device or the component of the device is in contact with a surface of a user that is sensitive to touch (e.g., a finger, palm, or other part of a user’s hand), the tactile output generated by the physical displacement will be interpreted

by the user as a tactile sensation corresponding to a perceived change in physical characteristics of the device or the component of the device. For example, movement of a touch-sensitive surface (e.g., a touch-sensitive display or trackpad) is, optionally, interpreted by the user as a “down click” or “up click” of a physical actuator button. In some cases, a user will feel a tactile sensation such as an “down click” or “up click” even when there is no movement of a physical actuator button associated with the touch-sensitive surface that is physically pressed (e.g., displaced) by the user’s movements. As another example, movement of the touch-sensitive surface is, optionally, interpreted or sensed by the user as “roughness” of the touch-sensitive surface, even when there is no change in smoothness of the touch-sensitive surface. While such interpretations of touch by a user will be subject to the individualized sensory perceptions of the user, there are many sensory perceptions of touch that are common to a large majority of users. Thus, when a tactile output is described as corresponding to a particular sensory perception of a user (e.g., an “up click,” a “down click,” “roughness”), unless otherwise stated, the generated tactile output corresponds to physical displacement of the device or a component thereof that will generate the described sensory perception for a typical (or average) user.

[0050] It will be appreciated that device **200** is only one example of a portable multifunction device, and that device **200** optionally has more or fewer components than shown, optionally combines two or more components, or optionally has a different configuration or arrangement of the components. The various components shown in FIG. 2A are implemented in hardware, software, or a combination of both hardware and software, including one or more signal processing and/or application-specific integrated circuits.

[0051] Memory **202** may include one or more computer-readable storage mediums. The computer-readable storage mediums may be tangible and non-transitory. Memory **202** may include high-speed random access memory and may also include non-volatile memory, such as one or more magnetic disk storage devices, flash memory devices, or other non-volatile solid-state memory devices. Memory controller **222** may control access to memory **202** by other components of device **200**.

[0052] In some examples, a non-transitory computer-readable storage medium of memory **202** can be used to store instructions (e.g., for performing aspects of processes **1000** and **1100**, described below) for use by or in connection with an instruction execution system, apparatus, or device, such as a computer-based system, processor-containing system, or other system that can fetch the instructions from the instruction execution system, apparatus, or device and execute the instructions. In other examples, the instructions (e.g., for performing aspects of processes **1000** and **1100**, described below) can be stored on a non-transitory computer-readable storage medium (not shown) of the server system **108** or can be divided between the non-transitory computer-readable storage medium of memory **202** and the non-transitory computer-readable storage medium of server system **108**. In the context of this document, a “non-transitory computer-readable storage medium” can be any medium that can contain or store the program for use by or in connection with the instruction execution system, apparatus, or device.

[0053] Peripherals interface **218** can be used to couple input and output peripherals of the device to CPU **220** and

memory **202**. The one or more processors **220** run or execute various software programs and/or sets of instructions stored in memory **202** to perform various functions for device **200** and to process data. In some embodiments, peripherals interface **218**, CPU **220**, and memory controller **222** may be implemented on a single chip, such as chip **204**. In some other embodiments, they may be implemented on separate chips.

[0054] RF (radio frequency) circuitry **208** receives and sends RF signals, also called electromagnetic signals. RF circuitry **208** converts electrical signals to/from electromagnetic signals and communicates with communications networks and other communications devices via the electromagnetic signals. RF circuitry **208** optionally includes well-known circuitry for performing these functions, including but not limited to an antenna system, an RF transceiver, one or more amplifiers, a tuner, one or more oscillators, a digital signal processor, a CODEC chipset, a subscriber identity module (SIM) card, memory, and so forth. RF circuitry **208** optionally communicates with networks, such as the Internet, also referred to as the World Wide Web (WWW), an intranet and/or a wireless network, such as a cellular telephone network, a wireless local area network (LAN) and/or a metropolitan area network (MAN), and other devices by wireless communication. The RF circuitry **208** optionally includes well-known circuitry for detecting near field communication (NFC) fields, such as by a short-range communication radio. The wireless communication optionally uses any of a plurality of communications standards, protocols, and technologies, including but not limited to Global System for Mobile Communications (GSM), Enhanced Data GSM Environment (EDGE), high-speed downlink packet access (HSDPA), high-speed uplink packet access (HSUPA), Evolution, Data-Only (EV-DO), HSPA, HSPA+, Dual-Cell HSPA (DC-HSPDA), long term evolution (LTE), near field communication (NFC), wideband code division multiple access (W-CDMA), code division multiple access (CDMA), time division multiple access (TDMA), Bluetooth, Bluetooth Low Energy (BTLE), Wireless Fidelity (Wi-Fi) (e.g., IEEE 802.11a, IEEE 802.11b, IEEE 802.11g, IEEE 802.11n, and/or IEEE 802.11ac), voice over Internet Protocol (VoIP), Wi-MAX, a protocol for e mail (e.g., Internet message access protocol (IMAP) and/or post office protocol (POP)), instant messaging (e.g., extensible messaging and presence protocol (XMPP), Session Initiation Protocol for Instant Messaging and Presence Leveraging Extensions (SIMPLE), Instant Messaging and Presence Service (IMPS)), and/or Short Message Service (SMS), or any other suitable communication protocol, including communication protocols not yet developed as of the filing date of this document.

[0055] Audio circuitry **210**, speaker **211**, and microphone **213** provide an audio interface between a user and device **200**. Audio circuitry **210** receives audio data from peripherals interface **218**, converts the audio data to an electrical signal, and transmits the electrical signal to speaker **211**. Speaker **211** converts the electrical signal to human-audible sound waves. Audio circuitry **210** also receives electrical signals converted by microphone **213** from sound waves. Audio circuitry **210** converts the electrical signal to audio data and transmits the audio data to peripherals interface **218** for processing. Audio data may be retrieved from and/or transmitted to memory **202** and/or RF circuitry **208** by peripherals interface **218**. In some embodiments, audio circuitry **210** also includes a headset jack (e.g., **312**, FIG. 3).

The headset jack provides an interface between audio circuitry **210** and removable audio input/output peripherals, such as output-only headphones or a headset with both output (e.g., a headphone for one or both ears) and input (e.g., a microphone).

[0056] I/O subsystem **206** couples input/output peripherals on device **200**, such as touch screen **212** and other input control devices **216**, to peripherals interface **218**. I/O subsystem **206** optionally includes display controller **256**, optical sensor controller **258**, intensity sensor controller **259**, haptic feedback controller **261**, and one or more input controllers **260** for other input or control devices. The one or more input controllers **260** receive/send electrical signals from/to other input control devices **216**. The other input control devices **216** optionally include physical buttons (e.g., push buttons, rocker buttons, etc.), dials, slider switches, joysticks, click wheels, and so forth. In some alternate embodiments, input controller(s) **260** are, optionally, coupled to any (or none) of the following: a keyboard, an infrared port, a USB port, and a pointer device such as a mouse. The one or more buttons (e.g., **308**, FIG. **3**) optionally include an up/down button for volume control of speaker **211** and/or microphone **213**. The one or more buttons optionally include a push button (e.g., **306**, FIG. **3**).

[0057] A quick press of the push button may disengage a lock of touch screen **212** or begin a process that uses gestures on the touch screen to unlock the device, as described in U.S. patent application Ser. No. 11/322,549, "Unlocking a Device by Performing Gestures on an Unlock Image," filed Dec. 23, 2005, U.S. Pat. No. 7,657,849, which is hereby incorporated by reference in its entirety. A longer press of the push button (e.g., **306**) may turn power to device **200** on or off. The user may be able to customize a functionality of one or more of the buttons. Touch screen **212** is used to implement virtual or soft buttons and one or more soft keyboards.

[0058] Touch-sensitive display **212** provides an input interface and an output interface between the device and a user. Display controller **256** receives and/or sends electrical signals from/to touch screen **212**. Touch screen **212** displays visual output to the user. The visual output may include graphics, text, icons, video, and any combination thereof (collectively termed "graphics"). In some embodiments, some or all of the visual output may correspond to user-interface objects.

[0059] Touch screen **212** has a touch-sensitive surface, sensor, or set of sensors that accepts input from the user based on haptic and/or tactile contact. Touch screen **212** and display controller **256** (along with any associated modules and/or sets of instructions in memory **202**) detect contact (and any movement or breaking of the contact) on touch screen **212** and convert the detected contact into interaction with user-interface objects (e.g., one or more soft keys, icons, web pages, or images) that are displayed on touch screen **212**. In an exemplary embodiment, a point of contact between touch screen **212** and the user corresponds to a finger of the user.

[0060] Touch screen **212** may use LCD (liquid crystal display) technology, LPD (light emitting polymer display) technology, or LED (light emitting diode) technology, although other display technologies may be used in other embodiments. Touch screen **212** and display controller **256** may detect contact and any movement or breaking thereof using any of a plurality of touch sensing technologies now

known or later developed, including but not limited to capacitive, resistive, infrared, and surface acoustic wave technologies, as well as other proximity sensor arrays or other elements for determining one or more points of contact with touch screen **212**. In an exemplary embodiment, projected mutual capacitance sensing technology is used, such as that found in the iPhone® and iPod Touch® from Apple Inc. of Cupertino, Calif.

[0061] A touch-sensitive display in some embodiments of touch screen **212** may be analogous to the multi-touch sensitive touchpads described in the following U.S. Pat. Nos. 6,323,846 (Westerman et al.), 6,570,557 (Westerman et al.), and/or 6,677,932 (Westerman), and/or U.S. Patent Publication 2002/0015024A1, each of which is hereby incorporated by reference in its entirety. However, touch screen **212** displays visual output from device **200**, whereas touch-sensitive touchpads do not provide visual output.

[0062] A touch-sensitive display in some embodiments of touch screen **212** may be as described in the following applications: (1) U.S. patent application Ser. No. 11/381,313, "Multipoint Touch Surface Controller," filed May 2, 2006; (2) U.S. patent application Ser. No. 10/840,862, "Multipoint Touchscreen," filed May 6, 2004; (3) U.S. patent application Ser. No. 10/903,964, "Gestures For Touch Sensitive Input Devices," filed Jul. 30, 2004; (4) U.S. patent application Ser. No. 11/048,264, "Gestures For Touch Sensitive Input Devices," filed Jan. 31, 2005; (5) U.S. patent application Ser. No. 11/038,590, "Mode-Based Graphical User Interfaces For Touch Sensitive Input Devices," filed Jan. 18, 2005; (6) U.S. patent application Ser. No. 11/228,758, "Virtual Input Device Placement On A Touch Screen User Interface," filed Sep. 16, 2005; (7) U.S. patent application Ser. No. 11/228,700, "Operation Of A Computer With A Touch Screen Interface," filed Sep. 16, 2005; (8) U.S. patent application Ser. No. 11/228,737, "Activating Virtual Keys Of A Touch-Screen Virtual Keyboard," filed Sep. 16, 2005; and (9) U.S. patent application Ser. No. 11/367,749, "Multi-Functional Hand-Held Device," filed Mar. 3, 2006. All of these applications are incorporated by reference herein in their entirety.

[0063] Touch screen **212** may have a video resolution in excess of 100 dpi. In some embodiments, the touch screen has a video resolution of approximately 160 dpi. The user may make contact with touch screen **212** using any suitable object or appendage, such as a stylus, a finger, and so forth. In some embodiments, the user interface is designed to work primarily with finger-based contacts and gestures, which can be less precise than stylus-based input due to the larger area of contact of a finger on the touch screen. In some embodiments, the device translates the rough finger-based input into a precise pointer/cursor position or command for performing the actions desired by the user.

[0064] In some embodiments, in addition to the touch screen, device **200** may include a touchpad (not shown) for activating or deactivating particular functions. In some embodiments, the touchpad is a touch-sensitive area of the device that, unlike the touch screen, does not display visual output. The touchpad may be a touch-sensitive surface that is separate from touch screen **212** or an extension of the touch-sensitive surface formed by the touch screen.

[0065] Device **200** also includes power system **262** for powering the various components. Power system **262** may include a power management system, one or more power sources (e.g., battery, alternating current (AC)), a recharging

system, a power failure detection circuit, a power converter or inverter, a power status indicator (e.g., a light-emitting diode (LED)) and any other components associated with the generation, management and distribution of power in portable devices.

[0066] Device **200** may also include one or more optical sensors **264**. FIG. 2A shows an optical sensor coupled to optical sensor controller **258** in I/O subsystem **206**. Optical sensor **264** may include charge-coupled device (CCD) or complementary metal-oxide semiconductor (CMOS) phototransistors. Optical sensor **264** receives light from the environment, projected through one or more lenses, and converts the light to data representing an image. In conjunction with imaging module **243** (also called a camera module), optical sensor **264** may capture still images or video. In some embodiments, an optical sensor is located on the back of device **200**, opposite touch screen display **212** on the front of the device so that the touch screen display may be used as a viewfinder for still and/or video image acquisition. In some embodiments, an optical sensor is located on the front of the device so that the user's image may be obtained for video conferencing while the user views the other video conference participants on the touch screen display. In some embodiments, the position of optical sensor **264** can be changed by the user (e.g., by rotating the lens and the sensor in the device housing) so that a single optical sensor **264** may be used along with the touch screen display for both video conferencing and still and/or video image acquisition.

[0067] Device **200** optionally also includes one or more contact intensity sensors **265**. FIG. 2A shows a contact intensity sensor coupled to intensity sensor controller **259** in I/O subsystem **206**. Contact intensity sensor **265** optionally includes one or more piezoresistive strain gauges, capacitive force sensors, electric force sensors, piezoelectric force sensors, optical force sensors, capacitive touch-sensitive surfaces, or other intensity sensors (e.g., sensors used to measure the force (or pressure) of a contact on a touch-sensitive surface). Contact intensity sensor **265** receives contact intensity information (e.g., pressure information or a proxy for pressure information) from the environment. In some embodiments, at least one contact intensity sensor is collocated with, or proximate to, a touch-sensitive surface (e.g., touch-sensitive display system **212**). In some embodiments, at least one contact intensity sensor is located on the back of device **200**, opposite touch screen display **212**, which is located on the front of device **200**.

[0068] Device **200** may also include one or more proximity sensors **266**. FIG. 2A shows proximity sensor **266** coupled to peripherals interface **218**. Alternately, proximity sensor **266** may be coupled to input controller **260** in I/O subsystem **206**. Proximity sensor **266** may perform as described in U.S. patent application Ser. Nos. 11/241,839, "Proximity Detector In Handheld Device"; 11/240,788, "Proximity Detector In Handheld Device"; 11/620,702, "Using Ambient Light Sensor To Augment Proximity Sensor Output"; 11/586,862, "Automated Response To And Sensing Of User Activity In Portable Devices"; and 11/638,251, "Methods And Systems For Automatic Configuration Of Peripherals," which are hereby incorporated by reference in their entirety. In some embodiments, the proximity sensor turns off and disables touch screen **212** when the multifunction device is placed near the user's ear (e.g., when the user is making a phone call).

[0069] Device **200** optionally also includes one or more tactile output generators **267**. FIG. 2A shows a tactile output generator coupled to haptic feedback controller **261** in I/O subsystem **206**. Tactile output generator **267** optionally includes one or more electroacoustic devices such as speakers or other audio components and/or electromechanical devices that convert energy into linear motion such as a motor, solenoid, electroactive polymer, piezoelectric actuator, electrostatic actuator, or other tactile output generating component (e.g., a component that converts electrical signals into tactile outputs on the device). Contact intensity sensor **265** receives tactile feedback generation instructions from haptic feedback module **233** and generates tactile outputs on device **200** that are capable of being sensed by a user of device **200**. In some embodiments, at least one tactile output generator is collocated with, or proximate to, a touch-sensitive surface (e.g., touch-sensitive display system **212**) and, optionally, generates a tactile output by moving the touch-sensitive surface vertically (e.g., in/out of a surface of device **200**) or laterally (e.g., back and forth in the same plane as a surface of device **200**). In some embodiments, at least one tactile output generator sensor is located on the back of device **200**, opposite touch screen display **212**, which is located on the front of device **200**.

[0070] Device **200** may also include one or more accelerometers **268**. FIG. 2A shows accelerometer **268** coupled to peripherals interface **218**. Alternately, accelerometer **268** may be coupled to an input controller **260** in I/O subsystem **206**. Accelerometer **268** may perform as described in U.S. Patent Publication No. 20050190059, "Acceleration-based Theft Detection System for Portable Electronic Devices," and U.S. Patent Publication No. 20060017692, "Methods And Apparatuses For Operating A Portable Device Based On An Accelerometer," both of which are incorporated by reference herein in their entirety. In some embodiments, information is displayed on the touch screen display in a portrait view or a landscape view based on an analysis of data received from the one or more accelerometers. Device **200** optionally includes, in addition to accelerometer(s) **268**, a magnetometer (not shown) and a GPS (or GLONASS or other global navigation system) receiver (not shown) for obtaining information concerning the location and orientation (e.g., portrait or landscape) of device **200**.

[0071] In some embodiments, the software components stored in memory **202** include operating system **226**, communication module (or set of instructions) **228**, contact/motion module (or set of instructions) **230**, graphics module (or set of instructions) **232**, text input module (or set of instructions) **234**, Global Positioning System (GPS) module (or set of instructions) **235**, Digital Assistant Client Module **229**, and applications (or sets of instructions) **236**. Further, memory **202** can store data and models, such as user data and models **231**. Furthermore, in some embodiments, memory **202** (FIG. 2A) or **470** (FIG. 4) stores device/global internal state **257**, as shown in FIGS. 2A and 4. Device/global internal state **257** includes one or more of: active application state, indicating which applications, if any, are currently active; display state, indicating what applications, views or other information occupy various regions of touch screen display **212**; sensor state, including information obtained from the device's various sensors and input control devices **216**; and location information concerning the device's location and/or attitude.

[0072] Operating system **226** (e.g., Darwin, RTXC, LINUX, UNIX, OS X, iOS, WINDOWS, or an embedded operating system such as VxWorks) includes various software components and/or drivers for controlling and managing general system tasks (e.g., memory management, storage device control, power management, etc.) and facilitates communication between various hardware and software components.

[0073] Communication module **228** facilitates communication with other devices over one or more external ports **224** and also includes various software components for handling data received by RF circuitry **208** and/or external port **224**. External port **224** (e.g., Universal Serial Bus (USB), FIREWIRE, etc.) is adapted for coupling directly to other devices or indirectly over a network (e.g., the Internet, wireless LAN, etc.). In some embodiments, the external port is a multi-pin (e.g., 30-pin) connector that is the same as, or similar to and/or compatible with, the 30-pin connector used on iPod® (trademark of Apple Inc.) devices.

[0074] Contact/motion module **230** optionally detects contact with touch screen **212** (in conjunction with display controller **256**) and other touch-sensitive devices (e.g., a touchpad or physical click wheel). Contact/motion module **230** includes various software components for performing various operations related to detection of contact, such as determining if contact has occurred (e.g., detecting a finger-down event), determining an intensity of the contact (e.g., the force or pressure of the contact or a substitute for the force or pressure of the contact), determining if there is movement of the contact and tracking the movement across the touch-sensitive surface (e.g., detecting one or more finger-dragging events), and determining if the contact has ceased (e.g., detecting a finger-up event or a break in contact). Contact/motion module **230** receives contact data from the touch-sensitive surface. Determining movement of the point of contact, which is represented by a series of contact data, optionally includes determining speed (magnitude), velocity (magnitude and direction), and/or an acceleration (a change in magnitude and/or direction) of the point of contact. These operations are, optionally, applied to single contacts (e.g., one finger contacts) or to multiple simultaneous contacts (e.g., “multitouch”/multiple finger contacts). In some embodiments, contact/motion module **230** and display controller **256** detect contact on a touchpad.

[0075] In some embodiments, contact/motion module **230** uses a set of one or more intensity thresholds to determine whether an operation has been performed by a user (e.g., to determine whether a user has “clicked” on an icon). In some embodiments, at least a subset of the intensity thresholds are determined in accordance with software parameters (e.g., the intensity thresholds are not determined by the activation thresholds of particular physical actuators and can be adjusted without changing the physical hardware of device **200**). For example, a mouse “click” threshold of a trackpad or touch screen display can be set to any of a large range of predefined threshold values without changing the trackpad or touch screen display hardware. Additionally, in some implementations, a user of the device is provided with software settings for adjusting one or more of the set of intensity thresholds (e.g., by adjusting individual intensity thresholds and/or by adjusting a plurality of intensity thresholds at once with a system-level click “intensity” parameter).

[0076] Contact/motion module **230** optionally detects a gesture input by a user. Different gestures on the touch-

sensitive surface have different contact patterns (e.g., different motions, timings, and/or intensities of detected contacts). Thus, a gesture is, optionally, detected by detecting a particular contact pattern. For example, detecting a finger tap gesture includes detecting a finger-down event followed by detecting a finger-up (liftoff) event at the same position (or substantially the same position) as the finger-down event (e.g., at the position of an icon). As another example, detecting a finger swipe gesture on the touch-sensitive surface includes detecting a finger-down event followed by detecting one or more finger-dragging events, and subsequently followed by detecting a finger-up (liftoff) event.

[0077] Graphics module **232** includes various known software components for rendering and displaying graphics on touch screen **212** or other display, including components for changing the visual impact (e.g., brightness, transparency, saturation, contrast, or other visual property) of graphics that are displayed. As used herein, the term “graphics” includes any object that can be displayed to a user, including, without limitation, text, web pages, icons (such as user-interface objects including soft keys), digital images, videos, animations, and the like.

[0078] In some embodiments, graphics module **232** stores data representing graphics to be used. Each graphic is, optionally, assigned a corresponding code. Graphics module **232** receives, from applications etc., one or more codes specifying graphics to be displayed along with, if necessary, coordinate data and other graphic property data, and then generates screen image data to output to display controller **256**.

[0079] Haptic feedback module **233** includes various software components for generating instructions used by tactile output generator(s) **267** to produce tactile outputs at one or more locations on device **200** in response to user interactions with device **200**.

[0080] Text input module **234**, which may be a component of graphics module **232**, provides soft keyboards for entering text in various applications (e.g., contacts **237**, e mail **240**, IM **241**, browser **247**, and any other application that needs text input).

[0081] GPS module **235** determines the location of the device and provides this information for use in various applications (e.g., to telephone **238** for use in location-based dialing; to camera **243** as picture/video metadata; and to applications that provide location-based services such as weather widgets, local yellow page widgets, and map/navigation widgets).

[0082] Digital assistant client module **229** can include various client-side digital assistant instructions to provide the client-side functionalities of the digital assistant. For example, digital assistant client module **229** can be capable of accepting voice input (e.g., speech input), text input, touch input, and/or gestural input through various user interfaces (e.g., microphone **213**, accelerometer(s) **268**, touch-sensitive display system **212**, optical sensor(s) **229**, other input control devices **216**, etc.) of portable multifunction device **200**. Digital assistant client module **229** can also be capable of providing output in audio (e.g., speech output), visual, and/or tactile forms through various output interfaces (e.g., speaker **211**, touch-sensitive display system **212**, tactile output generator(s) **267**, etc.) of portable multifunction device **200**. For example, output can be provided as voice, sound, alerts, text messages, menus, graphics, videos, animations, vibrations, and/or combinations of two or more of

the above. During operation, digital assistant client module 229 can communicate with DA server 106 using RF circuitry 208.

[0083] User data and models 231 can include various data associated with the user (e.g., user-specific vocabulary data, user preference data, user-specified name pronunciations, data from the user's electronic address book, to-do lists, shopping lists, etc.) to provide the client-side functionalities of the digital assistant. Further, user data and models 231 can include various models (e.g., speech recognition models, statistical language models, natural language processing models, ontology, task flow models, service models, etc.) for processing user input and determining user intent.

[0084] In some examples, digital assistant client module 229 can utilize the various sensors, subsystems, and peripheral devices of portable multifunction device 200 to gather additional information from the surrounding environment of the portable multifunction device 200 to establish a context associated with a user, the current user interaction, and/or the current user input. In some examples, digital assistant client module 229 can provide the contextual information or a subset thereof with the user input to DA server 106 to help infer the user's intent. In some examples, the digital assistant can also use the contextual information to determine how to prepare and deliver outputs to the user. Contextual information can be referred to as context data.

[0085] In some examples, the contextual information that accompanies the user input can include sensor information, e.g., lighting, ambient noise, ambient temperature, images or videos of the surrounding environment, etc. In some examples, the contextual information can also include the physical state of the device, e.g., device orientation, device location, device temperature, power level, speed, acceleration, motion patterns, cellular signals strength, etc. In some examples, information related to the software state of DA server 106, e.g., running processes, installed programs, past and present network activities, background services, error logs, resources usage, etc., and of portable multifunction device 200 can be provided to DA server 106 as contextual information associated with a user input.

[0086] In some examples, the digital assistant client module 229 can selectively provide information (e.g., user data 231) stored on the portable multifunction device 200 in response to requests from DA server 106. In some examples, digital assistant client module 229 can also elicit additional input from the user via a natural language dialogue or other user interfaces upon request by DA server 106. Digital assistant client module 229 can pass the additional input to DA server 106 to help DA server 106 in intent deduction and/or fulfillment of the user's intent expressed in the user request.

[0087] A more detailed description of a digital assistant is described below with reference to FIGS. 7A-C. It should be recognized that digital assistant client module 229 can include any number of the sub-modules of digital assistant module 726 described below.

[0088] Applications 236 may include the following modules (or sets of instructions), or a subset or superset thereof:

- [0089] Contacts module 237 (sometimes called an address book or contact list);
- [0090] Telephone module 238;
- [0091] Video conference module 239;
- [0092] Email client module 240;
- [0093] Instant messaging (IM) module 241;

- [0094] Workout support module 242;
- [0095] Camera module 243 for still and/or video images;
- [0096] Image management module 244;
- [0097] Video player module;
- [0098] Music player module;
- [0099] Browser module 247;
- [0100] Calendar module 248;
- [0101] Widget modules 249, which may include one or more of: weather widget 249-1, stocks widget 249-2, calculator widget 249-3, alarm clock widget 249-4, dictionary widget 249-5, and other widgets obtained by the user, as well as user-created widgets 249-6;
- [0102] Widget creator module 250 for making user-created widgets 249-6;
- [0103] Search module 251;
- [0104] Video and music player module 252, which merges video player module and music player module;
- [0105] Notes module 253;
- [0106] Map module 254; and/or
- [0107] Online video module 255.
- [0108] Examples of other applications 236 that may be stored in memory 202 include other word processing applications, other image editing applications, drawing applications, presentation applications, JAVA-enabled applications, encryption, digital rights management, voice recognition, and voice replication.
- [0109] In conjunction with touch screen 212, display controller 256, contact/motion module 230, graphics module 232, and text input module 234, contacts module 237 may be used to manage an address book or contact list (e.g., stored in application internal state 292 of contacts module 237 in memory 202 or memory 470), including: adding name(s) to the address book; deleting name(s) from the address book; associating telephone number(s), email address(es), physical address(es) or other information with a name; associating an image with a name; categorizing and sorting names; providing telephone numbers or email addresses to initiate and/or facilitate communications by telephone 238, video conference module 239, email 240, or IM 241; and so forth.
- [0110] In conjunction with RF circuitry 208, audio circuitry 210, speaker 211, microphone 213, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, and text input module 234, telephone module 238 may be used to enter a sequence of characters corresponding to a telephone number, access one or more telephone numbers in contacts module 237, modify a telephone number that has been entered, dial a respective telephone number, conduct a conversation, and disconnect or hang up when the conversation is completed. As noted above, the wireless communication may use any of a plurality of communications standards, protocols, and technologies.
- [0111] In conjunction with RF circuitry 208, audio circuitry 210, speaker 211, microphone 213, touch screen 212, display controller 256, optical sensor 264, optical sensor controller 258, contact/motion module 230, graphics module 232, text input module 234, contacts module 237, and telephone module 238, video conference module 239 includes executable instructions to initiate, conduct, and terminate a video conference between a user and one or more other participants in accordance with user instructions.
- [0112] In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, and text input module 234, email client module 240 includes executable instructions to create,

send, receive, and manage email in response to user instructions. In conjunction with image management module 244, email client module 240 makes it very easy to create and send emails with still or video images taken with camera module 243.

[0113] In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, and text input module 234, the instant messaging module 241 includes executable instructions to enter a sequence of characters corresponding to an instant message, to modify previously entered characters, to transmit a respective instant message (for example, using a Short Message Service (SMS) or Multimedia Message Service (MMS) protocol for telephony-based instant messages or using XMPP, SIMPLE, or IMPS for Internet-based instant messages), to receive instant messages, and to view received instant messages. In some embodiments, transmitted and/or received instant messages may include graphics, photos, audio files, video files and/or other attachments as are supported in an MMS and/or an Enhanced Messaging Service (EMS). As used herein, “instant messaging” refers to both telephony-based messages (e.g., messages sent using SMS or MMS) and Internet-based messages (e.g., messages sent using XMPP, SIMPLE, or IMPS).

[0114] In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, text input module 234, GPS module 235, map module 254, and music player module, workout support module 242 includes executable instructions to create workouts (e.g., with time, distance, and/or calorie burning goals); communicate with workout sensors (sports devices); receive workout sensor data; calibrate sensors used to monitor a workout; select and play music for a workout; and display, store, and transmit workout data.

[0115] In conjunction with touch screen 212, display controller 256, optical sensor(s) 264, optical sensor controller 258, contact/motion module 230, graphics module 232, and image management module 244, camera module 243 includes executable instructions to capture still images or video (including a video stream) and store them into memory 202, modify characteristics of a still image or video, or delete a still image or video from memory 202.

[0116] In conjunction with touch screen 212, display controller 256, contact/motion module 230, graphics module 232, text input module 234, and camera module 243, image management module 244 includes executable instructions to arrange, modify (e.g., edit), or otherwise manipulate, label, delete, present (e.g., in a digital slide show or album), and store still and/or video images.

[0117] In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, and text input module 234, browser module 247 includes executable instructions to browse the Internet in accordance with user instructions, including searching, linking to, receiving, and displaying web pages or portions thereof, as well as attachments and other files linked to web pages.

[0118] In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, text input module 234, email client module 240, and browser module 247, calendar module 248 includes executable instructions to create, display, modify,

and store calendars and data associated with calendars (e.g., calendar entries, to-do lists, etc.) in accordance with user instructions.

[0119] In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, text input module 234, and browser module 247, widget modules 249 are mini-applications that may be downloaded and used by a user (e.g., weather widget 249-1, stocks widget 249-2, calculator widget 249-3, alarm clock widget 249-4, and dictionary widget 249-5) or created by the user (e.g., user-created widget 249-6). In some embodiments, a widget includes an HTML (Hypertext Markup Language) file, a CSS (Cascading Style Sheets) file, and a JavaScript file. In some embodiments, a widget includes an XML (Extensible Markup Language) file and a JavaScript file (e.g., Yahoo! Widgets).

[0120] In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, text input module 234, and browser module 247, the widget creator module 250 may be used by a user to create widgets (e.g., turning a user-specified portion of a web page into a widget).

[0121] In conjunction with touch screen 212, display controller 256, contact/motion module 230, graphics module 232, and text input module 234, search module 251 includes executable instructions to search for text, music, sound, image, video, and/or other files in memory 202 that match one or more search criteria (e.g., one or more user-specified search terms) in accordance with user instructions.

[0122] In conjunction with touch screen 212, display controller 256, contact/motion module 230, graphics module 232, audio circuitry 210, speaker 211, RF circuitry 208, and browser module 247, video and music player module 252 includes executable instructions that allow the user to download and play back recorded music and other sound files stored in one or more file formats, such as MP3 or AAC files, and executable instructions to display, present, or otherwise play back videos (e.g., on touch screen 212 or on an external, connected display via external port 224). In some embodiments, device 200 optionally includes the functionality of an MP3 player, such as an iPod (trademark of Apple Inc.).

[0123] In conjunction with touch screen 212, display controller 256, contact/motion module 230, graphics module 232, and text input module 234, notes module 253 includes executable instructions to create and manage notes, to-do lists, and the like in accordance with user instructions.

[0124] In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, text input module 234, GPS module 235, and browser module 247, map module 254 may be used to receive, display, modify, and store maps and data associated with maps (e.g., driving directions, data on stores and other points of interest at or near a particular location, and other location-based data) in accordance with user instructions.

[0125] In conjunction with touch screen 212, display controller 256, contact/motion module 230, graphics module 232, audio circuitry 210, speaker 211, RF circuitry 208, text input module 234, email client module 240, and browser module 247, online video module 255 includes instructions that allow the user to access, browse, receive (e.g., by streaming and/or download), play back (e.g., on the touch screen or on an external, connected display via external port 224), send an email with a link to a particular online video,

and otherwise manage online videos in one or more file formats, such as H.264. In some embodiments, instant messaging module 241, rather than email client module 240, is used to send a link to a particular online video. Additional description of the online video application can be found in U.S. Provisional Patent Application No. 60/936,562, "Portable Multifunction Device, Method, and Graphical User Interface for Playing Online Videos," filed Jun. 20, 2007, and U.S. patent application Ser. No. 11/968,067, "Portable Multifunction Device, Method, and Graphical User Interface for Playing Online Videos," filed Dec. 31, 2007, the contents of which are hereby incorporated by reference in their entirety.

[0126] Each of the above-identified modules and applications corresponds to a set of executable instructions for performing one or more functions described above and the methods described in this application (e.g., the computer-implemented methods and other information processing methods described herein). These modules (e.g., sets of instructions) need not be implemented as separate software programs, procedures, or modules, and thus various subsets of these modules may be combined or otherwise rearranged in various embodiments. For example, video player module may be combined with music player module into a single module (e.g., video and music player module 252, FIG. 2A). In some embodiments, memory 202 may store a subset of the modules and data structures identified above. Furthermore, memory 202 may store additional modules and data structures not described above.

[0127] In some embodiments, device 200 is a device where operation of a predefined set of functions on the device is performed exclusively through a touch screen and/or a touchpad. By using a touch screen and/or a touchpad as the primary input control device for operation of device 200, the number of physical input control devices (such as push buttons, dials, and the like) on device 200 may be reduced.

[0128] The predefined set of functions that are performed exclusively through a touch screen and/or a touchpad optionally include navigation between user interfaces. In some embodiments, the touchpad, when touched by the user, navigates device 200 to a main, home, or root menu from any user interface that is displayed on device 200. In such embodiments, a "menu button" is implemented using a touchpad. In some other embodiments, the menu button is a physical push button or other physical input control device instead of a touchpad.

[0129] FIG. 2B is a block diagram illustrating exemplary components for event handling in accordance with some embodiments. In some embodiments, memory 202 (FIG. 2A) or 470 (FIG. 4) includes event sorter 270 (e.g., in operating system 226) and a respective application 236-1 (e.g., any of the aforementioned applications 237-251, 255, 480-490).

[0130] Event sorter 270 receives event information and determines the application 236-1 and application view 291 of application 236-1 to which to deliver the event information. Event sorter 270 includes event monitor 271 and event dispatcher module 274. In some embodiments, application 236-1 includes application internal state 292, which indicates the current application view(s) displayed on touch-sensitive display 212 when the application is active or executing. In some embodiments, device/global internal state 257 is used by event sorter 270 to determine which

application(s) is (are) currently active, and application internal state 292 is used by event sorter 270 to determine application views 291 to which to deliver event information.

[0131] In some embodiments, application internal state 292 includes additional information, such as one or more of: resume information to be used when application 236-1 resumes execution, user interface state information that indicates information being displayed or that is ready for display by application 236-1, a state queue for enabling the user to go back to a prior state or view of application 236-1, and a redo/undo queue of previous actions taken by the user.

[0132] Event monitor 271 receives event information from peripherals interface 218. Event information includes information about a sub-event (e.g., a user touch on touch-sensitive display 212, as part of a multi-touch gesture). Peripherals interface 218 transmits information it receives from I/O subsystem 206 or a sensor, such as proximity sensor 266, accelerometer(s) 268, and/or microphone 213 (through audio circuitry 210). Information that peripherals interface 218 receives from I/O subsystem 206 includes information from touch-sensitive display 212 or a touch-sensitive surface.

[0133] In some embodiments, event monitor 271 sends requests to the peripherals interface 218 at predetermined intervals. In response, peripherals interface 218 transmits event information. In other embodiments, peripherals interface 218 transmits event information only when there is a significant event (e.g., receiving an input above a predetermined noise threshold and/or for more than a predetermined duration).

[0134] In some embodiments, event sorter 270 also includes a hit view determination module 272 and/or an active event recognizer determination module 273.

[0135] Hit view determination module 272 provides software procedures for determining where a sub-event has taken place within one or more views when touch-sensitive display 212 displays more than one view. Views are made up of controls and other elements that a user can see on the display.

[0136] Another aspect of the user interface associated with an application is a set of views, sometimes herein called application views or user interface windows, in which information is displayed and touch-based gestures occur. The application views (of a respective application) in which a touch is detected may correspond to programmatic levels within a programmatic or view hierarchy of the application. For example, the lowest level view in which a touch is detected may be called the hit view, and the set of events that are recognized as proper inputs may be determined based, at least in part, on the hit view of the initial touch that begins a touch-based gesture.

[0137] Hit view determination module 272 receives information related to sub events of a touch-based gesture. When an application has multiple views organized in a hierarchy, hit view determination module 272 identifies a hit view as the lowest view in the hierarchy which should handle the sub-event. In most circumstances, the hit view is the lowest level view in which an initiating sub-event occurs (e.g., the first sub-event in the sequence of sub-events that form an event or potential event). Once the hit view is identified by the hit view determination module 272, the hit view typically receives all sub-events related to the same touch or input source for which it was identified as the hit view.

[0138] Active event recognizer determination module **273** determines which view or views within a view hierarchy should receive a particular sequence of sub-events. In some embodiments, active event recognizer determination module **273** determines that only the hit view should receive a particular sequence of sub-events. In other embodiments, active event recognizer determination module **273** determines that all views that include the physical location of a sub-event are actively involved views, and therefore determines that all actively involved views should receive a particular sequence of sub-events. In other embodiments, even if touch sub-events were entirely confined to the area associated with one particular view, views higher in the hierarchy would still remain as actively involved views.

[0139] Event dispatcher module **274** dispatches the event information to an event recognizer (e.g., event recognizer **280**). In embodiments including active event recognizer determination module **273**, event dispatcher module **274** delivers the event information to an event recognizer determined by active event recognizer determination module **273**. In some embodiments, event dispatcher module **274** stores in an event queue the event information, which is retrieved by a respective event receiver **282**.

[0140] In some embodiments, operating system **226** includes event sorter **270**. Alternatively, application **236-1** includes event sorter **270**. In yet other embodiments, event sorter **270** is a stand-alone module, or a part of another module stored in memory **202**, such as contact/motion module **230**.

[0141] In some embodiments, application **236-1** includes a plurality of event handlers **290** and one or more application views **291**, each of which includes instructions for handling touch events that occur within a respective view of the application's user interface. Each application view **291** of the application **236-1** includes one or more event recognizers **280**. Typically, a respective application view **291** includes a plurality of event recognizers **280**. In other embodiments, one or more of event recognizers **280** are part of a separate module, such as a user interface kit (not shown) or a higher level object from which application **236-1** inherits methods and other properties. In some embodiments, a respective event handler **290** includes one or more of: data updater **276**, object updater **277**, GUI updater **278**, and/or event data **279** received from event sorter **270**. Event handler **290** may utilize or call data updater **276**, object updater **277**, or GUI updater **278** to update the application internal state **292**. Alternatively, one or more of the application views **291** include one or more respective event handlers **290**. Also, in some embodiments, one or more of data updater **276**, object updater **277**, and GUI updater **278** are included in a respective application view **291**.

[0142] A respective event recognizer **280** receives event information (e.g., event data **279**) from event sorter **270** and identifies an event from the event information. Event recognizer **280** includes event receiver **282** and event comparator **284**. In some embodiments, event recognizer **280** also includes at least a subset of: metadata **283**, and event delivery instructions **288** (which may include sub-event delivery instructions).

[0143] Event receiver **282** receives event information from event sorter **270**. The event information includes information about a sub-event, for example, a touch or a touch movement. Depending on the sub-event, the event information also includes additional information, such as

location of the sub-event. When the sub-event concerns motion of a touch, the event information may also include speed and direction of the sub-event. In some embodiments, events include rotation of the device from one orientation to another (e.g., from a portrait orientation to a landscape orientation, or vice versa), and the event information includes corresponding information about the current orientation (also called device attitude) of the device.

[0144] Event comparator **284** compares the event information to predefined event or sub-event definitions and, based on the comparison, determines an event or sub event, or determines or updates the state of an event or sub-event. In some embodiments, event comparator **284** includes event definitions **286**. Event definitions **286** contain definitions of events (e.g., predefined sequences of sub-events), for example, event **1** (**287-1**), event **2** (**287-2**), and others. In some embodiments, sub-events in an event (**287**) include, for example, touch begin, touch end, touch movement, touch cancellation, and multiple touching. In one example, the definition for event **1** (**287-1**) is a double tap on a displayed object. The double tap, for example, comprises a first touch (touch begin) on the displayed object for a predetermined phase, a first liftoff (touch end) for a predetermined phase, a second touch (touch begin) on the displayed object for a predetermined phase, and a second liftoff (touch end) for a predetermined phase. In another example, the definition for event **2** (**287-2**) is a dragging on a displayed object. The dragging, for example, comprises a touch (or contact) on the displayed object for a predetermined phase, a movement of the touch across touch-sensitive display **212**, and liftoff of the touch (touch end). In some embodiments, the event also includes information for one or more associated event handlers **290**.

[0145] In some embodiments, event definition **287** includes a definition of an event for a respective user-interface object. In some embodiments, event comparator **284** performs a hit test to determine which user-interface object is associated with a sub-event. For example, in an application view in which three user-interface objects are displayed on touch-sensitive display **212**, when a touch is detected on touch-sensitive display **212**, event comparator **284** performs a hit test to determine which of the three user-interface objects is associated with the touch (sub-event). If each displayed object is associated with a respective event handler **290**, the event comparator uses the result of the hit test to determine which event handler **290** should be activated. For example, event comparator **284** selects an event handler associated with the sub-event and the object triggering the hit test.

[0146] In some embodiments, the definition for a respective event (**287**) also includes delayed actions that delay delivery of the event information until after it has been determined whether the sequence of sub-events does or does not correspond to the event recognizer's event type.

[0147] When a respective event recognizer **280** determines that the series of sub-events do not match any of the events in event definitions **286**, the respective event recognizer **280** enters an event impossible, event failed, or event ended state, after which it disregards subsequent sub-events of the touch-based gesture. In this situation, other event recognizers, if any, that remain active for the hit view continue to track and process sub-events of an ongoing touch-based gesture.

[0148] In some embodiments, a respective event recognizer 280 includes metadata 283 with configurable properties, flags, and/or lists that indicate how the event delivery system should perform sub-event delivery to actively involved event recognizers. In some embodiments, metadata 283 includes configurable properties, flags, and/or lists that indicate how event recognizers may interact, or are enabled to interact, with one another. In some embodiments, metadata 283 includes configurable properties, flags, and/or lists that indicate whether sub-events are delivered to varying levels in the view or programmatic hierarchy.

[0149] In some embodiments, a respective event recognizer 280 activates event handler 290 associated with an event when one or more particular sub-events of an event are recognized. In some embodiments, a respective event recognizer 280 delivers event information associated with the event to event handler 290. Activating an event handler 290 is distinct from sending (and deferred sending) sub-events to a respective hit view. In some embodiments, event recognizer 280 throws a flag associated with the recognized event, and event handler 290 associated with the flag catches the flag and performs a predefined process.

[0150] In some embodiments, event delivery instructions 288 include sub-event delivery instructions that deliver event information about a sub-event without activating an event handler. Instead, the sub-event delivery instructions deliver event information to event handlers associated with the series of sub-events or to actively involved views. Event handlers associated with the series of sub-events or with actively involved views receive the event information and perform a predetermined process.

[0151] In some embodiments, data updater 276 creates and updates data used in application 236-1. For example, data updater 276 updates the telephone number used in contacts module 237, or stores a video file used in video player module. In some embodiments, object updater 277 creates and updates objects used in application 236-1. For example, object updater 277 creates a new user-interface object or updates the position of a user-interface object. GUI updater 278 updates the GUI. For example, GUI updater 278 prepares display information and sends it to graphics module 232 for display on a touch-sensitive display.

[0152] In some embodiments, event handler(s) 290 includes or has access to data updater 276, object updater 277, and GUI updater 278. In some embodiments, data updater 276, object updater 277, and GUI updater 278 are included in a single module of a respective application 236-1 or application view 291. In other embodiments, they are included in two or more software modules.

[0153] It shall be understood that the foregoing discussion regarding event handling of user touches on touch-sensitive displays also applies to other forms of user inputs to operate multifunction devices 200 with input devices, not all of which are initiated on touch screens. For example, mouse movement and mouse button presses, optionally coordinated with single or multiple keyboard presses or holds; contact movements such as taps, drags, scrolls, etc. on touchpads; pen stylus inputs; movement of the device; oral instructions; detected eye movements; biometric inputs; and/or any combination thereof are optionally utilized as inputs corresponding to sub-events which define an event to be recognized.

[0154] FIG. 3 illustrates a portable multifunction device 200 having a touch screen 212 in accordance with some embodiments. The touch screen optionally displays one or

more graphics within user interface (UI) 300. In this embodiment, as well as others described below, a user is enabled to select one or more of the graphics by making a gesture on the graphics, for example, with one or more fingers 302 (not drawn to scale in the figure) or one or more styluses 303 (not drawn to scale in the figure). In some embodiments, selection of one or more graphics occurs when the user breaks contact with the one or more graphics. In some embodiments, the gesture optionally includes one or more taps, one or more swipes (from left to right, right to left, upward and/or downward), and/or a rolling of a finger (from right to left, left to right, upward and/or downward) that has made contact with device 200. In some implementations or circumstances, inadvertent contact with a graphic does not select the graphic. For example, a swipe gesture that sweeps over an application icon optionally does not select the corresponding application when the gesture corresponding to selection is a tap.

[0155] Device 200 may also include one or more physical buttons, such as “home” or menu button 304. As described previously, menu button 304 may be used to navigate to any application 236 in a set of applications that may be executed on device 200. Alternatively, in some embodiments, the menu button is implemented as a soft key in a GUI displayed on touch screen 212.

[0156] In one embodiment, device 200 includes touch screen 212, menu button 304, push button 306 for powering the device on/off and locking the device, volume adjustment button(s) 308, subscriber identity module (SIM) card slot 310, headset jack 312, and docking/charging external port 224. Push button 306 is, optionally, used to turn the power on/off on the device by depressing the button and holding the button in the depressed state for a predefined time interval; to lock the device by depressing the button and releasing the button before the predefined time interval has elapsed; and/or to unlock the device or initiate an unlock process. In an alternative embodiment, device 200 also accepts verbal input for activation or deactivation of some functions through microphone 213. Device 200 also, optionally, includes one or more contact intensity sensors 265 for detecting intensity of contacts on touch screen 212 and/or one or more tactile output generators 267 for generating tactile outputs for a user of device 200.

[0157] FIG. 4 is a block diagram of an exemplary multifunction device with a display and a touch-sensitive surface in accordance with some embodiments. Device 400 need not be portable. In some embodiments, device 400 is a laptop computer, a desktop computer, a tablet computer, a multimedia player device, a navigation device, an educational device (such as a child’s learning toy), a gaming system, or a control device (e.g., a home or industrial controller). Device 400 typically includes one or more processing units (CPUs) 410, one or more network or other communications interfaces 460, memory 470, and one or more communication buses 420 for interconnecting these components. Communication buses 420 optionally include circuitry (sometimes called a chipset) that interconnects and controls communications between system components. Device 400 includes input/output (I/O) interface 430 comprising display 440, which is typically a touch screen display. I/O interface 430 also optionally includes a keyboard and/or mouse (or other pointing device) 450 and touchpad 455, tactile output generator 457 for generating tactile outputs on device 400 (e.g., similar to tactile output generator(s) 267 described

above with reference to FIG. 2A), sensors 459 (e.g., optical, acceleration, proximity, touch-sensitive, and/or contact intensity sensors similar to contact intensity sensor(s) 265 described above with reference to FIG. 2A). Memory 470 includes high-speed random access memory, such as DRAM, SRAM, DDR RAM, or other random access solid state memory devices; and optionally includes non-volatile memory, such as one or more magnetic disk storage devices, optical disk storage devices, flash memory devices, or other non-volatile solid state storage devices. Memory 470 optionally includes one or more storage devices remotely located from CPU(s) 410. In some embodiments, memory 470 stores programs, modules, and data structures analogous to the programs, modules, and data structures stored in memory 202 of portable multifunction device 200 (FIG. 2A), or a subset thereof. Furthermore, memory 470 optionally stores additional programs, modules, and data structures not present in memory 202 of portable multifunction device 200. For example, memory 470 of device 400 optionally stores drawing module 480, presentation module 482, word processing module 484, website creation module 486, disk authoring module 488, and/or spreadsheet module 490, while memory 202 of portable multifunction device 200 (FIG. 2A) optionally does not store these modules.

[0158] Each of the above-identified elements in FIG. 4 may be stored in one or more of the previously mentioned memory devices. Each of the above-identified modules corresponds to a set of instructions for performing a function described above. The above-identified modules or programs (e.g., sets of instructions) need not be implemented as separate software programs, procedures, or modules, and thus various subsets of these modules may be combined or otherwise rearranged in various embodiments. In some embodiments, memory 470 may store a subset of the modules and data structures identified above. Furthermore, memory 470 may store additional modules and data structures not described above.

[0159] Attention is now directed towards embodiments of user interfaces that may be implemented on, for example, portable multifunction device 200.

[0160] FIG. 5A illustrates an exemplary user interface for a menu of applications on portable multifunction device 200 in accordance with some embodiments. Similar user interfaces may be implemented on device 400. In some embodiments, user interface 500 includes the following elements, or a subset or superset thereof:

[0161] Signal strength indicator(s) 502 for wireless communication(s), such as cellular and Wi-Fi signals;

[0162] Time 504;

[0163] Bluetooth indicator 505;

[0164] Battery status indicator 506;

[0165] Tray 508 with icons for frequently used applications, such as:

[0166] Icon 516 for telephone module 238, labeled "Phone," which optionally includes an indicator 514 of the number of missed calls or voicemail messages;

[0167] Icon 518 for email client module 240, labeled "Mail," which optionally includes an indicator 510 of the number of unread emails;

[0168] Icon 520 for browser module 247, labeled "Browser;" and

[0169] Icon 522 for video and music player module 252, also referred to as iPod (trademark of Apple Inc.) module 252, labeled "iPod;" and

[0170] Icons for other applications, such as:

[0171] Icon 524 for IM module 241, labeled "Messages;"

[0172] Icon 526 for calendar module 248, labeled "Calendar;"

[0173] Icon 528 for image management module 244, labeled "Photos;"

[0174] Icon 530 for camera module 243, labeled "Camera;"

[0175] Icon 532 for online video module 255, labeled "Online Video;"

[0176] Icon 534 for stocks widget 249-2, labeled "Stocks;"

[0177] Icon 536 for map module 254, labeled "Maps;"

[0178] Icon 538 for weather widget 249-1, labeled "Weather;"

[0179] Icon 540 for alarm clock widget 249-4, labeled "Clock;"

[0180] Icon 542 for workout support module 242, labeled "Workout Support;"

[0181] Icon 544 for notes module 253, labeled "Notes;" and

[0182] Icon 546 for a settings application or module, labeled "Settings," which provides access to settings for device 200 and its various applications 236.

[0183] It should be noted that the icon labels illustrated in FIG. 5A are merely exemplary. For example, icon 522 for video and music player module 252 may optionally be labeled "Music" or "Music Player." Other labels are, optionally, used for various application icons. In some embodiments, a label for a respective application icon includes a name of an application corresponding to the respective application icon. In some embodiments, a label for a particular application icon is distinct from a name of an application corresponding to the particular application icon.

[0184] FIG. 5B illustrates an exemplary user interface on a device (e.g., device 400, FIG. 4) with a touch-sensitive surface 551 (e.g., a tablet or touchpad 455, FIG. 4) that is separate from the display 550 (e.g., touch screen display 212). Device 400 also, optionally, includes one or more contact intensity sensors (e.g., one or more of sensors 457) for detecting intensity of contacts on touch-sensitive surface 551 and/or one or more tactile output generators 459 for generating tactile outputs for a user of device 400.

[0185] Although some of the examples which follow will be given with reference to inputs on touch screen display 212 (where the touch-sensitive surface and the display are combined), in some embodiments, the device detects inputs on a touch-sensitive surface that is separate from the display, as shown in FIG. 5B. In some embodiments, the touch-sensitive surface (e.g., 551 in FIG. 5B) has a primary axis (e.g., 552 in FIG. 5B) that corresponds to a primary axis (e.g., 553 in FIG. 5B) on the display (e.g., 550). In accordance with these embodiments, the device detects contacts (e.g., 560 and 562 in FIG. 5B) with the touch-sensitive surface 551 at locations that correspond to respective locations on the display (e.g., in FIG. 5B, 560 corresponds to 568 and 562 corresponds to 570). In this way, user inputs (e.g., contacts 560 and 562, and movements thereof) detected by the device on the touch-sensitive surface (e.g., 551 in FIG. 5B) are used by the device to manipulate the user interface on the display (e.g., 550 in FIG. 5B) of the multifunction device when the touch-sensitive surface is separate from the

display. It should be understood that similar methods are, optionally, used for other user interfaces described herein.

[0186] Additionally, while the following examples are given primarily with reference to finger inputs (e.g., finger contacts, finger tap gestures, finger swipe gestures), it should be understood that, in some embodiments, one or more of the finger inputs are replaced with input from another input device (e.g., a mouse-based input or stylus input). For example, a swipe gesture is, optionally, replaced with a mouse click (e.g., instead of a contact) followed by movement of the cursor along the path of the swipe (e.g., instead of movement of the contact). As another example, a tap gesture is, optionally, replaced with a mouse click while the cursor is located over the location of the tap gesture (e.g., instead of detection of the contact followed by ceasing to detect the contact). Similarly, when multiple user inputs are simultaneously detected, it should be understood that multiple computer mice are, optionally, used simultaneously, or a mouse and finger contacts are, optionally, used simultaneously.

[0187] FIG. 6A illustrates exemplary personal electronic device 600. Device 600 includes body 602. In some embodiments, device 600 can include some or all of the features described with respect to devices 200 and 400 (e.g., FIGS. 2A-4B). In some embodiments, device 600 has touch-sensitive display screen 604, hereafter touch screen 604. Alternatively, or in addition to touch screen 604, device 600 has a display and a touch-sensitive surface. As with devices 200 and 400, in some embodiments, touch screen 604 (or the touch-sensitive surface) may have one or more intensity sensors for detecting intensity of contacts (e.g., touches) being applied. The one or more intensity sensors of touch screen 604 (or the touch-sensitive surface) can provide output data that represents the intensity of touches. The user interface of device 600 can respond to touches based on their intensity, meaning that touches of different intensities can invoke different user interface operations on device 600.

[0188] Techniques for detecting and processing touch intensity may be found, for example, in related applications: International Patent Application Serial No. PCT/US2013/040061, titled "Device, Method, and Graphical User Interface for Displaying User Interface Objects Corresponding to an Application," filed May 8, 2013, and International Patent Application Serial No. PCT/US2013/069483, titled "Device, Method, and Graphical User Interface for Transitioning Between Touch Input to Display Output Relationships," filed Nov. 11, 2013, each of which is hereby incorporated by reference in their entirety.

[0189] In some embodiments, device 600 has one or more input mechanisms 606 and 608. Input mechanisms 606 and 608, if included, can be physical. Examples of physical input mechanisms include push buttons and rotatable mechanisms. In some embodiments, device 600 has one or more attachment mechanisms. Such attachment mechanisms, if included, can permit attachment of device 600 with, for example, hats, eyewear, earrings, necklaces, shirts, jackets, bracelets, watch straps, chains, trousers, belts, shoes, purses, backpacks, and so forth. These attachment mechanisms may permit device 600 to be worn by a user.

[0190] FIG. 6B depicts exemplary personal electronic device 600. In some embodiments, device 600 can include some or all of the components described with respect to FIGS. 2A, 2B, and 4. Device 600 has bus 612 that operatively couples I/O section 614 with one or more computer

processors 616 and memory 618. I/O section 614 can be connected to display 604, which can have touch-sensitive component 622 and, optionally, touch-intensity sensitive component 624. In addition, I/O section 614 can be connected with communication unit 630 for receiving application and operating system data, using Wi-Fi, Bluetooth, near field communication (NFC), cellular, and/or other wireless communication techniques. Device 600 can include input mechanisms 606 and/or 608. Input mechanism 606 may be a rotatable input device or a depressible and rotatable input device, for example. Input mechanism 608 may be a button, in some examples.

[0191] Input mechanism 608 may be a microphone, in some examples. Personal electronic device 600 can include various sensors, such as GPS sensor 632, accelerometer 634, directional sensor 640 (e.g., compass), gyroscope 636, motion sensor 638, and/or a combination thereof, all of which can be operatively connected to I/O section 614.

[0192] Memory 618 of personal electronic device 600 can include one or more non-transitory computer-readable storage mediums, for storing computer-executable instructions, which, when executed by one or more computer processors 616, for example, can cause the computer processors to perform the techniques described below, including processes 1000 and 1100 (FIGS. 10-11). Personal electronic device 600 is not limited to the components and configuration of FIG. 6B, but can include other or additional components in multiple configurations.

[0193] As used here, the term "affordance" refers to a user-interactive graphical user interface object that may be displayed on the display screen of devices 200, 400, and/or 600 (FIGS. 2, 4, and 6). For example, an image (e.g., icon), a button, and text (e.g., link) may each constitute an affordance.

[0194] As used herein, the term "focus selector" refers to an input element that indicates a current part of a user interface with which a user is interacting. In some implementations that include a cursor or other location marker, the cursor acts as a "focus selector" so that when an input (e.g., a press input) is detected on a touch-sensitive surface (e.g., touchpad 455 in FIG. 4 or touch-sensitive surface 551 in FIG. 5B) while the cursor is over a particular user interface element (e.g., a button, window, slider or other user interface element), the particular user interface element is adjusted in accordance with the detected input. In some implementations that include a touch screen display (e.g., touch-sensitive display system 212 in FIG. 2A or touch screen 212 in FIG. 5A) that enables direct interaction with user interface elements on the touch screen display, a detected contact on the touch screen acts as a "focus selector" so that when an input (e.g., a press input by the contact) is detected on the touch screen display at a location of a particular user interface element (e.g., a button, window, slider, or other user interface element), the particular user interface element is adjusted in accordance with the detected input. In some implementations, focus is moved from one region of a user interface to another region of the user interface without corresponding movement of a cursor or movement of a contact on a touch screen display (e.g., by using a tab key or arrow keys to move focus from one button to another button); in these implementations, the focus selector moves in accordance with movement of focus between different regions of the user interface. Without regard to the specific form taken by the focus selector, the focus selector is

generally the user interface element (or contact on a touch screen display) that is controlled by the user so as to communicate the user's intended interaction with the user interface (e.g., by indicating, to the device, the element of the user interface with which the user is intending to interact). For example, the location of a focus selector (e.g., a cursor, a contact, or a selection box) over a respective button while a press input is detected on the touch-sensitive surface (e.g., a touchpad or touch screen) will indicate that the user is intending to activate the respective button (as opposed to other user interface elements shown on a display of the device).

[0195] As used in the specification and claims, the term "characteristic intensity" of a contact refers to a characteristic of the contact based on one or more intensities of the contact. In some embodiments, the characteristic intensity is based on multiple intensity samples. The characteristic intensity is, optionally, based on a predefined number of intensity samples, or a set of intensity samples collected during a predetermined time period (e.g., 0.05, 0.1, 0.2, 0.5, 1, 2, 5, 10 seconds) relative to a predefined event (e.g., after detecting the contact, prior to detecting liftoff of the contact, before or after detecting a start of movement of the contact, prior to detecting an end of the contact, before or after detecting an increase in intensity of the contact, and/or before or after detecting a decrease in intensity of the contact). A characteristic intensity of a contact is, optionally based on one or more of: a maximum value of the intensities of the contact, a mean value of the intensities of the contact, an average value of the intensities of the contact, a top 10 percentile value of the intensities of the contact, a value at the half maximum of the intensities of the contact, a value at the 90 percent maximum of the intensities of the contact, or the like. In some embodiments, the duration of the contact is used in determining the characteristic intensity (e.g., when the characteristic intensity is an average of the intensity of the contact over time). In some embodiments, the characteristic intensity is compared to a set of one or more intensity thresholds to determine whether an operation has been performed by a user. For example, the set of one or more intensity thresholds may include a first intensity threshold and a second intensity threshold. In this example, a contact with a characteristic intensity that does not exceed the first threshold results in a first operation, a contact with a characteristic intensity that exceeds the first intensity threshold and does not exceed the second intensity threshold results in a second operation, and a contact with a characteristic intensity that exceeds the second threshold results in a third operation. In some embodiments, a comparison between the characteristic intensity and one or more thresholds is used to determine whether or not to perform one or more operations (e.g., whether to perform a respective operation or forgo performing the respective operation) rather than being used to determine whether to perform a first operation or a second operation.

[0196] In some embodiments, a portion of a gesture is identified for purposes of determining a characteristic intensity. For example, a touch-sensitive surface may receive a continuous swipe contact transitioning from a start location and reaching an end location, at which point the intensity of the contact increases. In this example, the characteristic intensity of the contact at the end location may be based on only a portion of the continuous swipe contact, and not the entire swipe contact (e.g., only the portion of the swipe

contact at the end location). In some embodiments, a smoothing algorithm may be applied to the intensities of the swipe contact prior to determining the characteristic intensity of the contact. For example, the smoothing algorithm optionally includes one or more of: an unweighted sliding-average smoothing algorithm, a triangular smoothing algorithm, a median filter smoothing algorithm, and/or an exponential smoothing algorithm. In some circumstances, these smoothing algorithms eliminate narrow spikes or dips in the intensities of the swipe contact for purposes of determining a characteristic intensity.

[0197] The intensity of a contact on the touch-sensitive surface may be characterized relative to one or more intensity thresholds, such as a contact-detection intensity threshold, a light press intensity threshold, a deep press intensity threshold, and/or one or more other intensity thresholds. In some embodiments, the light press intensity threshold corresponds to an intensity at which the device will perform operations typically associated with clicking a button of a physical mouse or a trackpad. In some embodiments, the deep press intensity threshold corresponds to an intensity at which the device will perform operations that are different from operations typically associated with clicking a button of a physical mouse or a trackpad. In some embodiments, when a contact is detected with a characteristic intensity below the light press intensity threshold (e.g., and above a nominal contact-detection intensity threshold below which the contact is no longer detected), the device will move a focus selector in accordance with movement of the contact on the touch-sensitive surface without performing an operation associated with the light press intensity threshold or the deep press intensity threshold. Generally, unless otherwise stated, these intensity thresholds are consistent between different sets of user interface figures.

[0198] An increase of characteristic intensity of the contact from an intensity below the light press intensity threshold to an intensity between the light press intensity threshold and the deep press intensity threshold is sometimes referred to as a "light press" input. An increase of characteristic intensity of the contact from an intensity below the deep press intensity threshold to an intensity above the deep press intensity threshold is sometimes referred to as a "deep press" input. An increase of characteristic intensity of the contact from an intensity below the contact-detection intensity threshold to an intensity between the contact-detection intensity threshold and the light press intensity threshold is sometimes referred to as detecting the contact on the touch-surface. A decrease of characteristic intensity of the contact from an intensity above the contact-detection intensity threshold to an intensity below the contact-detection intensity threshold is sometimes referred to as detecting liftoff of the contact from the touch-surface. In some embodiments, the contact-detection intensity threshold is zero. In some embodiments, the contact-detection intensity threshold is greater than zero.

[0199] In some embodiments described herein, one or more operations are performed in response to detecting a gesture that includes a respective press input or in response to detecting the respective press input performed with a respective contact (or a plurality of contacts), where the respective press input is detected based at least in part on detecting an increase in intensity of the contact (or plurality of contacts) above a press-input intensity threshold. In some embodiments, the respective operation is performed in

response to detecting the increase in intensity of the respective contact above the press-input intensity threshold (e.g., a “down stroke” of the respective press input). In some embodiments, the press input includes an increase in intensity of the respective contact above the press-input intensity threshold and a subsequent decrease in intensity of the contact below the press-input intensity threshold, and the respective operation is performed in response to detecting the subsequent decrease in intensity of the respective contact below the press-input threshold (e.g., an “up stroke” of the respective press input).

[0200] In some embodiments, the device employs intensity hysteresis to avoid accidental inputs sometimes termed “jitter,” where the device defines or selects a hysteresis intensity threshold with a predefined relationship to the press-input intensity threshold (e.g., the hysteresis intensity threshold is X intensity units lower than the press-input intensity threshold or the hysteresis intensity threshold is 75%, 90%, or some reasonable proportion of the press-input intensity threshold). Thus, in some embodiments, the press input includes an increase in intensity of the respective contact above the press-input intensity threshold and a subsequent decrease in intensity of the contact below the hysteresis intensity threshold that corresponds to the press-input intensity threshold, and the respective operation is performed in response to detecting the subsequent decrease in intensity of the respective contact below the hysteresis intensity threshold (e.g., an “up stroke” of the respective press input). Similarly, in some embodiments, the press input is detected only when the device detects an increase in intensity of the contact from an intensity at or below the hysteresis intensity threshold to an intensity at or above the press-input intensity threshold and, optionally, a subsequent decrease in intensity of the contact to an intensity at or below the hysteresis intensity, and the respective operation is performed in response to detecting the press input (e.g., the increase in intensity of the contact or the decrease in intensity of the contact, depending on the circumstances).

[0201] For ease of explanation, the descriptions of operations performed in response to a press input associated with a press-input intensity threshold or in response to a gesture including the press input are, optionally, triggered in response to detecting either: an increase in intensity of a contact above the press-input intensity threshold, an increase in intensity of a contact from an intensity below the hysteresis intensity threshold to an intensity above the press-input intensity threshold, a decrease in intensity of the contact below the press-input intensity threshold, and/or a decrease in intensity of the contact below the hysteresis intensity threshold corresponding to the press-input intensity threshold. Additionally, in examples where an operation is described as being performed in response to detecting a decrease in intensity of a contact below the press-input intensity threshold, the operation is, optionally, performed in response to detecting a decrease in intensity of the contact below a hysteresis intensity threshold corresponding to, and lower than, the press-input intensity threshold.

3. Digital Assistant System

[0202] FIG. 7A illustrates a block diagram of digital assistant system 700 in accordance with various examples. In some examples, digital assistant system 700 can be implemented on a standalone computer system. In some examples, digital assistant system 700 can be distributed

across multiple computers. In some examples, some of the modules and functions of the digital assistant can be divided into a server portion and a client portion, where the client portion resides on one or more user devices (e.g., devices 104, 122, 200, 400, or 600) and communicates with the server portion (e.g., server system 108) through one or more networks, e.g., as shown in FIG. 1. In some examples, digital assistant system 700 can be an implementation of server system 108 (and/or DA server 106) shown in FIG. 1. It should be noted that digital assistant system 700 is only one example of a digital assistant system, and that digital assistant system 700 can have more or fewer components than shown, may combine two or more components, or may have a different configuration or arrangement of the components. The various components shown in FIG. 7A can be implemented in hardware, software instructions for execution by one or more processors, firmware, including one or more signal processing and/or application specific integrated circuits, or a combination thereof.

[0203] Digital assistant system 700 can include memory 702, one or more processors 704, input/output (I/O) interface 706, and network communications interface 708. These components can communicate with one another over one or more communication buses or signal lines 710.

[0204] In some examples, memory 702 can include a non-transitory computer-readable medium, such as high-speed random access memory and/or a non-volatile computer-readable storage medium (e.g., one or more magnetic disk storage devices, flash memory devices, or other non-volatile solid-state memory devices).

[0205] In some examples, I/O interface 706 can couple input/output devices 716 of digital assistant system 700, such as displays, keyboards, touch screens, and microphones, to user interface module 722. I/O interface 706, in conjunction with user interface module 722, can receive user inputs (e.g., voice input, keyboard inputs, touch inputs, etc.) and processes them accordingly. In some examples, e.g., when the digital assistant is implemented on a standalone user device, digital assistant system 700 can include any of the components and I/O communication interfaces described with respect to devices 200, 400, or 600 in FIGS. 2A, 4, 6A-B, respectively. In some examples, digital assistant system 700 can represent the server portion of a digital assistant implementation, and can interact with the user through a client-side portion residing on a user device (e.g., devices 104, 200, 400, or 600).

[0206] In some examples, the network communications interface 708 can include wired communication port(s) 712 and/or wireless transmission and reception circuitry 714. The wired communication port(s) can receive and send communication signals via one or more wired interfaces, e.g., Ethernet, Universal Serial Bus (USB), FIREWIRE, etc. The wireless circuitry 714 can receive and send RF signals and/or optical signals from/to communications networks and other communications devices. The wireless communications can use any of a plurality of communications standards, protocols, and technologies, such as GSM, EDGE, CDMA, TDMA, Bluetooth, Wi-Fi, VoIP, Wi-MAX, or any other suitable communication protocol. Network communications interface 708 can enable communication between digital assistant system 700 with networks, such as the Internet, an intranet, and/or a wireless network, such as a

cellular telephone network, a wireless local area network (LAN), and/or a metropolitan area network (MAN), and other devices.

[0207] In some examples, memory 702, or the computer-readable storage media of memory 702, can store programs, modules, instructions, and data structures including all or a subset of: operating system 718, communications module 720, user interface module 722, one or more applications 724, and digital assistant module 726. In particular, memory 702, or the computer-readable storage media of memory 702, can store instructions for performing processes 1000 and 1100, described below. One or more processors 704 can execute these programs, modules, and instructions, and reads/writes from/to the data structures.

[0208] Operating system 718 (e.g., Darwin, RTXC, LINUX, UNIX, iOS, OS X, WINDOWS, or an embedded operating system such as VxWorks) can include various software components and/or drivers for controlling and managing general system tasks (e.g., memory management, storage device control, power management, etc.) and facilitates communications between various hardware, firmware, and software components.

[0209] Communications module 720 can facilitate communications between digital assistant system 700 with other devices over network communications interface 708. For example, communications module 720 can communicate with RF circuitry 208 of electronic devices such as devices 200, 400, and 600 shown in FIG. 2A, 4, 6A-B, respectively. Communications module 720 can also include various components for handling data received by wireless circuitry 714 and/or wired communications port 712.

[0210] User interface module 722 can receive commands and/or inputs from a user via I/O interface 706 (e.g., from a keyboard, touch screen, pointing device, controller, and/or microphone), and generate user interface objects on a display. User interface module 722 can also prepare and deliver outputs (e.g., speech, sound, animation, text, icons, vibrations, haptic feedback, light, etc.) to the user via the I/O interface 706 (e.g., through displays, audio channels, speakers, touch-pads, etc.).

[0211] Applications 724 can include programs and/or modules that are configured to be executed by one or more processors 704. For example, if the digital assistant system is implemented on a standalone user device, applications 724 can include user applications, such as games, a calendar application, a navigation application, or an email application. If digital assistant system 700 is implemented on a server, applications 724 can include resource management applications, diagnostic applications, or scheduling applications, for example.

[0212] Memory 702 can also store digital assistant module 726 (or the server portion of a digital assistant). In some examples, digital assistant module 726 can include the following sub-modules, or a subset or superset thereof: input/output processing module 728, speech-to-text (STT) processing module 730, natural language processing module 732, dialogue flow processing module 734, task flow processing module 736, service processing module 738, and speech synthesis module 740. Each of these modules can have access to one or more of the following systems or data and models of the digital assistant module 726, or a subset or superset thereof: ontology 760, vocabulary index 744, user data 748, task flow models 754, service models 756, and ASR systems.

[0213] In some examples, using the processing modules, data, and models implemented in digital assistant module 726, the digital assistant can perform at least some of the following: converting speech input into text; identifying a user's intent expressed in a natural language input received from the user; actively eliciting and obtaining information needed to fully infer the user's intent (e.g., by disambiguating words, games, intentions, etc.); determining the task flow for fulfilling the inferred intent; and executing the task flow to fulfill the inferred intent.

[0214] In some examples, as shown in FIG. 7B, I/O processing module 728 can interact with the user through I/O devices 716 in FIG. 7A or with a user device (e.g., devices 104, 200, 400, or 600) through network communications interface 708 in FIG. 7A to obtain user input (e.g., a speech input) and to provide responses (e.g., as speech outputs) to the user input. I/O processing module 728 can optionally obtain contextual information associated with the user input from the user device, along with or shortly after the receipt of the user input. The contextual information can include user-specific data, vocabulary, and/or preferences relevant to the user input. In some examples, the contextual information also includes software and hardware states of the user device at the time the user request is received, and/or information related to the surrounding environment of the user at the time that the user request was received. In some examples, I/O processing module 728 can also send follow-up questions to, and receive answers from, the user regarding the user request. When a user request is received by I/O processing module 728 and the user request can include speech input, I/O processing module 728 can forward the speech input to STT processing module 730 (or speech recognizer) for speech-to-text conversions.

[0215] STT processing module 730 can include one or more ASR systems. The one or more ASR systems can process the speech input that is received through I/O processing module 728 to produce a recognition result. Each ASR system can include a front-end speech pre-processor. The front-end speech pre-processor can extract representative features from the speech input. For example, the front-end speech pre-processor can perform a Fourier transform on the speech input to extract spectral features that characterize the speech input as a sequence of representative multi-dimensional vectors. Further, each ASR system can include one or more speech recognition models (e.g., acoustic models and/or language models) and can implement one or more speech recognition engines. Examples of speech recognition models can include Hidden Markov Models, Gaussian-Mixture Models, Deep Neural Network Models, n-gram language models, and other statistical models. Examples of speech recognition engines can include the dynamic time warping based engines and weighted finite-state transducers (WFST) based engines. The one or more speech recognition models and the one or more speech recognition engines can be used to process the extracted representative features of the front-end speech pre-processor to produce intermediate recognitions results (e.g., phonemes, phonemic strings, and sub-words), and ultimately, text recognition results (e.g., words, word strings, or sequence of tokens). In some examples, the speech input can be processed at least partially by a third-party service or on the user's device (e.g., device 104, 200, 400, or 600) to produce the recognition result. Once STT processing module 730 produces recognition results containing a text string

(e.g., words, or sequence of words, or sequence of tokens), the recognition result can be passed to natural language processing module 732 for intent deduction.

[0216] More details on the speech-to-text processing are described in U.S. Utility application Ser. No. 13/236,942 for “Consolidating Speech Recognition Results,” filed on Sep. 20, 2011, the entire disclosure of which is incorporated herein by reference.

[0217] In some examples, STT processing module 730 can include and/or access a vocabulary of recognizable words via phonetic alphabet conversion module 731. Each vocabulary word can be associated with one or more candidate pronunciations of the word represented in a speech recognition phonetic alphabet. In particular, the vocabulary of recognizable words can include a word that is associated with a plurality of candidate pronunciations. For example, the vocabulary may include the word “tomato” that is associated with the candidate pronunciations of /təˈmeɪroʊ/ and /təˈmatoʊ/. Further, vocabulary words can be associated with custom candidate pronunciations that are based on previous speech inputs from the user. Such custom candidate pronunciations can be stored in STT processing module 730 and can be associated with a particular user via the user’s profile on the device. In some examples, the candidate pronunciations for words can be determined based on the spelling of the word and one or more linguistic and/or phonetic rules. In some examples, the candidate pronunciations can be manually generated, e.g., based on known canonical pronunciations.

[0218] In some examples, the candidate pronunciations can be ranked based on the commonness of the candidate pronunciation. For example, the candidate pronunciation /təˈmeɪroʊ/ can be ranked higher than /təˈmatoʊ/, because the former is a more commonly used pronunciation (e.g., among all users, for users in a particular geographical region, or for any other appropriate subset of users). In some examples, candidate pronunciations can be ranked based on whether the candidate pronunciation is a custom candidate pronunciation associated with the user. For example, custom candidate pronunciations can be ranked higher than canonical candidate pronunciations. This can be useful for recognizing proper nouns having a unique pronunciation that deviates from canonical pronunciation. In some examples, candidate pronunciations can be associated with one or more speech characteristics, such as geographic origin, nationality, or ethnicity. For example, the candidate pronunciation /təˈmeɪroʊ/ can be associated with the United States, whereas the candidate pronunciation /təˈmatoʊ/ can be associated with Great Britain. Further, the rank of the candidate pronunciation can be based on one or more characteristics (e.g., geographic origin, nationality, ethnicity, etc.) of the user stored in the user’s profile on the device. For example, it can be determined from the user’s profile that the user is associated with the United States. Based on the user being associated with the United States, the candidate pronunciation /təˈmeɪroʊ/ (associated with the United States) can be ranked higher than the candidate pronunciation /təˈmatoʊ/ (associated with Great Britain). In some examples, one of the ranked candidate pronunciations can be selected as a predicted pronunciation (e.g., the most likely pronunciation).

[0219] When a speech input is received, STT processing module 730 can be used to determine the phonemes corresponding to the speech input (e.g., using an acoustic model), and then attempt to determine words that match the phonemes (e.g., using a language model). For example, if STT processing module 730 can first identify the sequence of phonemes /təˈmeɪroʊ/ corresponding to a portion of the speech input, it can then determine, based on vocabulary index 744, that this sequence corresponds to the word “tomato.”

[0220] In some examples, STT processing module 730 can use approximate matching techniques to determine words in an utterance. Thus, for example, the STT processing module 730 can determine that the sequence of phonemes /təˈmeɪroʊ/ corresponds to the word “tomato,” even if that particular sequence of phonemes is not one of the candidate sequence of phonemes for that word.

[0221] Natural language processing module 732 (“natural language processor”) of the digital assistant can take the sequence of words or tokens (“token sequence”) generated by STT processing module 730, and attempt to associate the token sequence with one or more “actionable intents” recognized by the digital assistant. An “actionable intent” can represent a task that can be performed by the digital assistant, and can have an associated task flow implemented in task flow models 754. The associated task flow can be a series of programmed actions and steps that the digital assistant takes in order to perform the task. The scope of a digital assistant’s capabilities can be dependent on the number and variety of task flows that have been implemented and stored in task flow models 754, or in other words, on the number and variety of “actionable intents” that the digital assistant recognizes. The effectiveness of the digital assistant, however, can also be dependent on the assistant’s ability to infer the correct “actionable intent(s)” from the user request expressed in natural language.

[0222] In some examples, in addition to the sequence of words or tokens obtained from STT processing module 730, natural language processing module 732 can also receive contextual information associated with the user request, e.g., from I/O processing module 728. The natural language processing module 732 can optionally use the contextual information to clarify, supplement, and/or further define the information contained in the token sequence received from STT processing module 730. The contextual information can include, for example, user preferences, hardware, and/or software states of the user device, sensor information collected before, during, or shortly after the user request, prior interactions (e.g., dialogue) between the digital assistant and the user, and the like. As described herein, contextual information can be dynamic, and can change with time, location, content of the dialogue, and other factors.

[0223] In some examples, the natural language processing can be based on, e.g., ontology 760. Ontology 760 can be a hierarchical structure containing many nodes, each node representing either an “actionable intent” or a “property” relevant to one or more of the “actionable intents” or other “properties.” As noted above, an “actionable intent” can represent a task that the digital assistant is capable of performing, i.e., it is “actionable” or can be acted on. A “property” can represent a parameter associated with an actionable intent or a sub-aspect of another property. A linkage between an actionable intent node and a property node in ontology 760 can define how a parameter repre-

sented by the property node pertains to the task represented by the actionable intent node.

[0224] In some examples, ontology **760** can be made up of actionable intent nodes and property nodes. Within ontology **760**, each actionable intent node can be linked to one or more property nodes either directly or through one or more intermediate property nodes. Similarly, each property node can be linked to one or more actionable intent nodes either directly or through one or more intermediate property nodes. For example, as shown in FIG. 7C, ontology **760** can include a “restaurant reservation” node (i.e., an actionable intent node). Property nodes “restaurant,” “date/time” (for the reservation), and “party size” can each be directly linked to the actionable intent node (i.e., the “restaurant reservation” node).

[0225] In addition, property nodes “cuisine,” “price range,” “phone number,” and “location” can be sub-nodes of the property node “restaurant,” and can each be linked to the “restaurant reservation” node (i.e., the actionable intent node) through the intermediate property node “restaurant.” For another example, as shown in FIG. 7C, ontology **760** can also include a “set reminder” node (i.e., another actionable intent node). Property nodes “date/time” (for setting the reminder) and “subject” (for the reminder) can each be linked to the “set reminder” node. Since the property “date/time” can be relevant to both the task of making a restaurant reservation and the task of setting a reminder, the property node “date/time” can be linked to both the “restaurant reservation” node and the “set reminder” node in ontology **760**.

[0226] An actionable intent node, along with its linked concept nodes, can be described as a “domain.” In the present discussion, each domain can be associated with a respective actionable intent, and refers to the group of nodes (and the relationships there between) associated with the particular actionable intent. For example, ontology **760** shown in FIG. 7C can include an example of restaurant reservation domain **762** and an example of reminder domain **764** within ontology **760**. The restaurant reservation domain includes the actionable intent node “restaurant reservation,” property nodes “restaurant,” “date/time,” and “party size,” and sub-property nodes “cuisine,” “price range,” “phone number,” and “location.” Reminder domain **764** can include the actionable intent node “set reminder,” and property nodes “subject” and “date/time.” In some examples, ontology **760** can be made up of many domains. Each domain can share one or more property nodes with one or more other domains. For example, the “date/time” property node can be associated with many different domains (e.g., a scheduling domain, a travel reservation domain, a movie ticket domain, etc.), in addition to restaurant reservation domain **762** and reminder domain **764**.

[0227] While FIG. 7C illustrates two example domains within ontology **760**, other domains can include, for example, “find a movie,” “initiate a phone call,” “find directions,” “schedule a meeting,” “send a message,” and “provide an answer to a question,” “read a list,” “providing navigation instructions,” “provide instructions for a task” and so on. A “send a message” domain can be associated with a “send a message” actionable intent node, and may further include property nodes such as “recipient(s),” “message type,” and “message body.” The property node “recipient” can be further defined, for example, by the sub-property nodes such as “recipient name” and “message address.”

[0228] In some examples, ontology **760** can include all the domains (and hence actionable intents) that the digital assistant is capable of understanding and acting upon. In some examples, ontology **760** can be modified, such as by adding or removing entire domains or nodes, or by modifying relationships between the nodes within the ontology **760**.

[0229] In some examples, nodes associated with multiple related actionable intents can be clustered under a “super domain” in ontology **760**. For example, a “travel” super-domain can include a cluster of property nodes and actionable intent nodes related to travel. The actionable intent nodes related to travel can include “airline reservation,” “hotel reservation,” “car rental,” “get directions,” “find points of interest,” and so on. The actionable intent nodes under the same super domain (e.g., the “travel” super domain) can have many property nodes in common. For example, the actionable intent nodes for “airline reservation,” “hotel reservation,” “car rental,” “get directions,” and “find points of interest” can share one or more of the property nodes “start location,” “destination,” “departure date/time,” “arrival date/time,” and “party size.”

[0230] In some examples, each node in ontology **760** can be associated with a set of words and/or phrases that are relevant to the property or actionable intent represented by the node. The respective set of words and/or phrases associated with each node can be the so-called “vocabulary” associated with the node. The respective set of words and/or phrases associated with each node can be stored in vocabulary index **744** in association with the property or actionable intent represented by the node. For example, returning to FIG. 7B, the vocabulary associated with the node for the property of “restaurant” can include words such as “food,” “drinks,” “cuisine,” “hungry,” “eat,” “pizza,” “fast food,” “meal,” and so on. For another example, the vocabulary associated with the node for the actionable intent of “initiate a phone call” can include words and phrases such as “call,” “phone,” “dial,” “ring,” “call this number,” “make a call to,” and so on. The vocabulary index **744** can optionally include words and phrases in different languages.

[0231] Natural language processing module **732** can receive the token sequence (e.g., a text string) from STT processing module **730**, and determine what nodes are implicated by the words in the token sequence. In some examples, if a word or phrase in the token sequence is found to be associated with one or more nodes in ontology **760** (via vocabulary index **744**), the word or phrase can “trigger” or “activate” those nodes. Based on the quantity and/or relative importance of the activated nodes, natural language processing module **732** can select one of the actionable intents as the task that the user intended the digital assistant to perform. In some examples, the domain that has the most “triggered” nodes can be selected. In some examples, the domain having the highest confidence value (e.g., based on the relative importance of its various triggered nodes) can be selected. In some examples, the domain can be selected based on a combination of the number and the importance of the triggered nodes. In some examples, additional factors are considered in selecting the node as well, such as whether the digital assistant has previously correctly interpreted a similar request from a user.

[0232] User data **748** can include user-specific information, such as user-specific vocabulary, user preferences, user address, user’s default and secondary languages, user’s

contact list, and other short-term or long-term information for each user. In some examples, natural language processing module 732 can use the user-specific information to supplement the information contained in the user input to further define the user intent. For example, for a user request “invite my friends to my birthday party,” natural language processing module 732 can be able to access user data 748 to determine who the “friends” are and when and where the “birthday party” would be held, rather than requiring the user to provide such information explicitly in his/her request.

[0233] Other details of searching an ontology based on a token string is described in U.S. Utility application Ser. No. 12/341,743 for “Method and Apparatus for Searching Using An Active Ontology,” filed Dec. 22, 2008, the entire disclosure of which is incorporated herein by reference.

[0234] In some examples, once natural language processing module 732 identifies an actionable intent (or domain) based on the user request, natural language processing module 732 can generate a structured query to represent the identified actionable intent. In some examples, the structured query can include parameters for one or more nodes within the domain for the actionable intent, and at least some of the parameters are populated with the specific information and requirements specified in the user request. For example, the user may say “Make me a dinner reservation at a sushi place at 7.” In this case, natural language processing module 732 can be able to correctly identify the actionable intent to be “restaurant reservation” based on the user input. According to the ontology, a structured query for a “restaurant reservation” domain may include parameters such as {Cuisine}, {Time}, {Date}, {Party Size}, and the like. In some examples, based on the speech input and the text derived from the speech input using STT processing module 730, natural language processing module 732 can generate a partial structured query for the restaurant reservation domain, where the partial structured query includes the parameters {Cuisine 32 “Sushi”} and {Time=“7 pm”}. However, in this example, the user’s utterance contains insufficient information to complete the structured query associated with the domain. Therefore, other necessary parameters such as {Party Size} and {Date} may not be specified in the structured query based on the information currently available. In some examples, natural language processing module 732 can populate some parameters of the structured query with received contextual information. For example, in some examples, if the user requested a sushi restaurant “near me,” natural language processing module 732 can populate a {location} parameter in the structured query with GPS coordinates from the user device.

[0235] In some examples, natural language processing module 732 can pass the generated structured query (including any completed parameters) to task flow processing module 736 (“task flow processor”). Task flow processing module 736 can be configured to receive the structured query from natural language processing module 732, complete the structured query, if necessary, and perform the actions required to “complete” the user’s ultimate request. In some examples, the various procedures necessary to complete these tasks can be provided in task flow models 754. In some examples, task flow models 754 can include procedures for obtaining additional information from the user and task flows for performing actions associated with the actionable intent.

[0236] As described above, in order to complete a structured query, task flow processing module 736 may need to initiate additional dialogue with the user in order to obtain additional information, and/or disambiguate potentially ambiguous utterances. When such interactions are necessary, task flow processing module 736 can invoke dialogue flow processing module 734 to engage in a dialogue with the user. In some examples, dialogue flow processing module 734 can determine how (and/or when) to ask the user for the additional information and receives and processes the user responses. The questions can be provided to and answers can be received from the users through I/O processing module 728. In some examples, dialogue flow processing module 734 can present dialogue output to the user via audio and/or visual output, and receives input from the user via spoken or physical (e.g., clicking) responses. Continuing with the example above, when task flow processing module 736 invokes dialogue flow processing module 734 to determine the “party size” and “date” information for the structured query associated with the domain “restaurant reservation,” dialogue flow processing module 734 can generate questions such as “For how many people?” and “On which day?” to pass to the user. Once answers are received from the user, dialogue flow processing module 734 can then populate the structured query with the missing information, or pass the information to task flow processing module 736 to complete the missing information from the structured query.

[0237] Once task flow processing module 736 has completed the structured query for an actionable intent, task flow processing module 736 can proceed to perform the ultimate task associated with the actionable intent. Accordingly, task flow processing module 736 can execute the steps and instructions in the task flow model according to the specific parameters contained in the structured query. For example, the task flow model for the actionable intent of “restaurant reservation” can include steps and instructions for contacting a restaurant and actually requesting a reservation for a particular party size at a particular time. For example, using a structured query such as: {restaurant reservation, restaurant=ABC Café, date=3/12/2012, time=7 pm, party size=5}, task flow processing module 736 can perform the steps of: (1) logging onto a server of the ABC Café or a restaurant reservation system such as OPENTABLE®, (2) entering the date, time, and party size information in a form on the website, (3) submitting the form, and (4) making a calendar entry for the reservation in the user’s calendar.

[0238] In some examples, task flow processing module 736 can employ the assistance of service processing module 738 (“service processing module”) to complete a task requested in the user input or to provide an informational answer requested in the user input. For example, service processing module 738 can act on behalf of task flow processing module 736 to make a phone call, set a calendar entry, invoke a map search, invoke or interact with other user applications installed on the user device, and invoke or interact with third-party services (e.g., a restaurant reservation portal, a social networking website, a banking portal, etc.). In some examples, the protocols and application programming interfaces (API) required by each service can be specified by a respective service model among service models 756. Service processing module 738 can access the appropriate service model for a service and generate requests for the service in accordance with the protocols and APIs required by the service according to the service model.

[0239] For example, if a restaurant has enabled an online reservation service, the restaurant can submit a service model specifying the necessary parameters for making a reservation and the APIs for communicating the values of the necessary parameter to the online reservation service. When requested by task flow processing module 736, service processing module 738 can establish a network connection with the online reservation service using the web address stored in the service model, and send the necessary parameters of the reservation (e.g., time, date, party size) to the online reservation interface in a format according to the API of the online reservation service.

[0240] In some examples, natural language processing module 732, dialogue flow processing module 734, and task flow processing module 736 can be used collectively and iteratively to infer and define the user's intent, obtain information to further clarify and refine the user intent, and finally generate a response (i.e., an output to the user, or the completion of a task) to fulfill the user's intent. The generated response can be a dialogue response to the speech input that at least partially fulfills the user's intent. Further, in some examples, the generated response can be output as a speech output. In these examples, the generated response can be sent to speech synthesis module 740 (e.g., speech synthesizer) where it can be processed to synthesize the dialogue response in speech form. In yet other examples, the generated response can be data content relevant to satisfying a user request in the speech input.

[0241] Speech synthesis module 740 can be configured to synthesize speech outputs for presentation to the user. Speech synthesis module 740 synthesizes speech outputs based on text provided by the digital assistant. For example, the generated dialogue response can be in the form of a text string. Speech synthesis module 740 can convert the text string to an audible speech output. Speech synthesis module 740 can use any appropriate speech synthesis technique in order to generate speech outputs from text, including, but not limited to, concatenative synthesis, unit selection synthesis, diphone synthesis, domain-specific synthesis, formant synthesis, articulatory synthesis, hidden Markov model (HMM) based synthesis, and sinewave synthesis. In some examples, speech synthesis module 740 can be configured to synthesize individual words based on phonemic strings corresponding to the words. For example, a phonemic string can be associated with a word in the generated dialogue response. The phonemic string can be stored in metadata associated with the word. Speech synthesis model 740 can be configured to directly process the phonemic string in the metadata to synthesize the word in speech form.

[0242] In some examples, instead of (or in addition to) using speech synthesis module 740, speech synthesis can be performed on a remote device (e.g., the server system 108), and the synthesized speech can be sent to the user device for output to the user. For example, this can occur in some implementations where outputs for a digital assistant are generated at a server system. And because server systems generally have more processing power or resources than a user device, it can be possible to obtain higher quality speech outputs than would be practical with client-side synthesis.

[0243] Additional details on digital assistants can be found in the U.S. Utility application Ser. No. 12/987,982, entitled "Intelligent Automated Assistant," filed Jan. 10, 2011, U.S. Utility application Ser. No. 13/251,088, entitled "Generating

and Processing Task Items That Represent Tasks to Perform," filed Sep. 30, 2011, U.S. Utility application Ser. No. 15/144,618, entitled "Intelligent Device Identification," filed May 2, 2016, and U.S. Utility application Ser. No. 14/503,105, entitled "Intelligent Assistant for Home Automation," filed Sep. 30, 2014, the entire disclosures of which are incorporated herein by reference.

4. Exemplary Techniques for Intelligent Device Arbitration

[0244] FIGS. 8A-C illustrate exemplary techniques for intelligent device arbitration in accordance with some embodiments. These figures are also used to illustrate the processes described below, including the processes in FIGS. 10A-C.

[0245] FIG. 8A illustrates electronic devices 802, 804, 806, and 808 of user 800. One or more of the devices 802-808 may be any of devices 104, 122, 200, 400, 600, and 1200 (FIGS. 1, 2A, 3, 4, 5A, 6A-6B, and 12) in some embodiments. In some examples, electronic device 802 is a wearable electronic device, such as a smart watch, and is, optionally, powered off when in a lowered position, as illustrated. In some examples, electronic device 804 is a mobile device, such as a mobile phone; electronic device 806 is an electronic device having relatively large display capabilities, such as a television; electronic device 808 is a device having auditory output capabilities, such as a speaker dock. While the illustrated example is described herein with reference to electronic devices 802-808, it should be appreciated that a fewer or greater number of devices may be used in other implementations. Further, electronic devices 802-808 may be associated with different users (not shown).

[0246] Each of the electronic devices 802-808 is capable of sampling audio inputs from user 800. For example, electronic device 804 samples audio inputs through its microphone to listen for spoken instructions from user 800. In some examples, one or more of electronic devices 802-808 continuously sample audio input. In some examples, one or more of electronic devices 802-808 begin to sample audio input in response to detecting proximity of user 800. In some examples, one or more of electronic devices 802-808 begin to sample audio input in response to a particular input by user 800, such as a spoken trigger. FIG. 8A illustrates user 800 providing spoken instruction 834, "Hey Siri find me a TV episode". In one embodiment, each of electronic devices 802-808 samples audio input and detects spoken instruction 834, respectively.

[0247] In response to detecting spoken instruction 834, electronic devices 802-808 initiate an arbitration process to identify (e.g., determine) an electronic device for responding to the spoken instruction 834 from user 800. By way of example, each of electronic devices 802-808 broadcasts a set of values based on the audio input as sampled on the respective device. Each set of values may include one or more values. For example, electronic device 804 (implemented as a mobile phone) broadcasts a first set of one or more values based on the spoken instruction 834 as sampled on electronic device 804, while electronic device 808 (implemented as a speaker dock) broadcasts a second set of one or more values based on the spoken instruction 834 as sampled on electronic device 808. Electronic devices 802 and 806 broadcast respective sets of one or more values as well. The sets of one or more values can be broadcasted by electronic devices 802-808 via any unidirectional broadcast communications standards, protocols, and/or technologies

known now or in the future. In some examples, one or more sets of one or more values are broadcasted using Bluetooth Low Energy (BTLE) advertising mode. Other communication methods can be used, including but not limited to WiFi (e.g., any 802.11 compliant communication), NFC, infrared, sound wave, etc., or any combination thereof.

[0248] In some examples, a set of one or more values broadcasted by an electronic device includes any number of values. One exemplary value indicates an energy level of the audio input as sampled on the electronic device. The energy level of the sampled audio input may, for instance, be indicative of the proximity of the electronic device to the user. In some examples, the energy level is measured using known metrics for audio quality, such as signal-to-noise ratio, sound pressure, or a combination thereof.

[0249] Another exemplary value indicates an acoustic fingerprint of the audio input, that is, whether the audio input is likely provided by a particular user. In some examples, the electronic device analyzes the audio input and calculates a confidence value that expresses the likelihood that the audio input originates from the particular user (e.g., a user that has enrolled in the virtual assistant service on the electronic device). In some examples, the electronic device determines whether the audio input is from an authorized user by comparing the confidence value with a predetermined threshold value, and broadcasts a value based on the comparison. In some examples, the set of one or more values may include a value corresponding to the confidence value as well.

[0250] Another exemplary value indicates a type of the electronic device. For example, as illustrated in FIG. 8A, the electronic device **804** is a mobile phone. It will be appreciated that predetermined values can be used by electronic devices **802-808** to indicate different device types, including but not limited to, speaker, television, smart watch, laptop, tablet, mobile device, set-top box, headphones, or any combination thereof. In some examples, each of the electronic devices **802-808** broadcasts a “device type” value only when the spoken instruction specifies a task to be performed, as discussed in more details below.

[0251] Another exemplary value indicates a state of the electronic device. In some examples, the value can, for instance, be indicative of whether the electronic device has been in an active state (e.g., having received a user input) within a predetermined amount of time before receiving the audio input. In at least one example, the value can indicate whether the electronic device is in a locked state. In some examples, the value can also indicate whether the user has recently activated the device or the virtual assistant service on the device. In some examples, the value can convey details regarding the recent user input, such as a time stamp, a type of input (e.g., a physical touch, a raising gesture), etc.

[0252] Another exemplary value indicates location information of an electronic device. In some examples, the value indicates a geographic location of the electronic device, for example, using GPS coordinates of the electronic device, or indicates a named location of the electronic device (e.g., the user’s living room).

[0253] The above-discussed exemplary values correspond to various metrics relevant to implementing intelligent device arbitration as described herein. It will be appreciated that any number of values corresponding to these metrics can be broadcasted and/or that some or all of the values can be used to provide a single value using one or more

functions. The single value may thereafter be broadcasted during device arbitration as described herein. It will be further appreciated that the one or more functions can assign different weights to different values, respectively.

[0254] In some examples, the electronic device determines respective values and/or broadcasts the values in response to detecting a spoken trigger. With reference to FIGS. 8A-C, the spoken trigger is the phrase “Hey Siri”. In response to the spoken trigger “Hey Siri”, each of the electronic devices **802-808** broadcasts the respective set of one or more values, as described. If an audio input does not contain a spoken trigger, the electronic device foregoes determining and/or broadcasting a set of one or more values.

[0255] Each of the electronic devices **802-808** may receive sets of values from other devices. By way of example, the electronic device **802** may receive sets of values from electronic devices **804-808**, the electronic device **804** may receive sets of values from electronic devices **802, 806, 808**, and so on. After sets of values have been exchanged among electronic devices **802-808**, each device determines whether it is to respond to the audio input by analyzing the sets of values, as discussed below.

[0256] In some examples, each device determines whether to respond to the spoken instruction **834** based on the energy level values broadcasted by electronic devices **802-808**. Each of the electronic devices **802-808** compares a respective “energy level” value with the “energy level” values broadcasted by other electronic devices. In some examples, an electronic device responds to the audio input if the electronic device has broadcasted the highest “energy level” value. As discussed above, a higher “energy level” value can be indicative of greater proximity to the user. Because a device closest to a user may be associated with a highest energy level value, it is beneficial to have an electronic device that has broadcasted the highest “energy level” value respond to the spoken instruction **934**.

[0257] In some examples, each device determines whether to respond to the spoken instruction **834** based on the state values broadcasted by electronic devices **802-808**. Each of the electronic devices **802-808** compares a respective “state” value with the “state” values broadcasted by other electronic devices. As discussed above, the “state” value(s) of an electronic device can encompass information regarding the operation of the electronic device, such as whether it is locked, whether it has been recently activated, whether it has recently received particular user inputs. In some examples, an electronic device responds to the audio input if the user has recently provided an input indicative of an intent to interact with the particular electronic device. Exemplary inputs include an active gesture (e.g. on a wearable electronic device), a physical touch (e.g. on a touch-sensitive screen of a device), etc. In some examples, an electronic device responds to the audio input if, based on the broadcasted “state” values, the electronic device is the only device that is in an unlocked state.

[0258] In some examples, each device determines whether to respond to the spoken instruction **834** based on the acoustic fingerprint values broadcasted by electronic devices **802-808**. Each of the electronic devices **802-808** compares a respective “acoustic fingerprint” value with the values broadcasted by other electronic devices. In some examples, an electronic device responds to the audio input if, based on the broadcasted values, it is the only electronic device that has recognized the user who provided the audio input as an

authorized user. For example, if user **800** has enrolled in the virtual assistant service only on electronic device **802** (implemented as a smart watch), electronic device **802** responds to the audio inputs from user **800**.

[0259] In some examples, each of the electronic devices **802-808** determines whether to respond to the spoken instruction based on the device type values broadcasted by electronic devices **802-808**. Each of the electronic devices **802-808** compares a respective “device type” value(s) with the values broadcasted by other electronic devices. Device type is particularly relevant to intelligent device arbitration when the user’s input specifies a task to be performed, such as “Hey Siri, find me a TV episode” or “Hey Siri, drive me there”. In some examples, the electronic device is to respond to the audio input if, based on the broadcasted values, the electronic device is the only device of a type that can handle the task specified in the audio input.

[0260] In some examples, the electronic device obtains a list of acceptable device types that can handle the specified task (e.g., by locally analyzing the audio input and/or by receiving the list from one or more servers), and determines whether the electronic device is to respond to the audio input based on the broadcasted “device type” values and the list of acceptable device types. In some examples, one or more servers receive data representing the sampled audio input from one or more of electronic devices **802-808**, derive a user intent based on the data, and identify a task having one or more parameters based on the user intent. In some examples, once the one or more servers determine the task based on the audio input, the one or more servers transmit the task, parameters (if any), and a list of acceptable device types for handling the task to the electronic devices that have sampled the audio input (e.g., electronic devices **802-808**). Additional details regarding the identification of a task from a natural language input can be found, for example, in U.S. Utility application Ser. No. 15/144,618, entitled “Intelligent Device Identification”, filed May 2, 2016, which is hereby incorporated by reference in its entirety.

[0261] It will be appreciated that the arbitration processes described herein are exemplary, and that, using one or more numerical and/or logical functions and algorithms, some or all of the values above can be factored, alone or in combination, into the determination of whether an electronic device is to respond to an audio input. In some examples, each of electronic devices **802-808** broadcasts multiple values (e.g., an aggregated score and a “device type” value). Accordingly, electronic devices **802-808** may arbitrate in accordance with one or more predetermined algorithms. In some examples, each of electronic devices **802-808** broadcasts a single score, which is calculated according to one or more functions using some or all of the above-discussed values. For example, a single score for broadcasting can be calculated based on a predetermined weighting of received audio quality (e.g., signal-to-noise ratio), type of device, ability to perform task, and device state. Further, it will be appreciated that a variety of logical and/or numerical functions and algorithms can be used by each of electronic devices **802-808**. It should be further appreciated that adjustments to the functions and algorithms can be implemented to adjust the prioritization of the above-described factors.

[0262] The electronic device can respond to an audio input by providing a visual output (e.g., a display notification or LED toggle), an auditory output, a haptic output, or a combination thereof, in some examples. For example, elec-

tronic device **804** can respond to the audio input with a visual output (e.g. displaying a transcript of the audio input) and an audio output (e.g., “Looking for episodes . . .”). If the electronic device determines not to respond to the audio input, the electronic device can forego responding to the audio input by entering an inactive mode (e.g. sleep mode).

[0263] With reference to FIG. 8A, each of electronic device **802** (implemented as a smart watch), electronic device **804** (implemented as a mobile phone), electronic device **808** (implemented as a speaker dock), and electronic device **806** (implemented as a television) generates and broadcasts a set of one or more values after sampling spoken instruction **834** from user **800**. In this example, each set of one or more values includes an “energy level” value and a “device type” value. By comparing the broadcasted “energy level” values, electronic device **804** determines that it has broadcasted the highest “energy level” value among electronic devices **802-808**, indicating that electronic device **804** is in relatively close proximity to user **800**.

[0264] However, electronic device **804** further analyzes the broadcasted “device type” values in light of the task specified in the spoken instruction **834**. As discussed above, spoken instruction **834** can be processed and resolved into one or more tasks locally or remotely. For example, electronic device **804** can resolve the spoken instruction **834** into a task locally, or receive the task along with a list of acceptable device types for handling the task from one or more servers. In this example, electronic device **804** determines that it is not of an acceptable (or preferred) device type for handling the task specified in the audio input (i.e., playback of a video). As such, electronic device **804** foregoes responding to the audio input.

[0265] In some examples, electronic device **804** further determines whether any of the rest of the electronic devices can handle the specified task. Electronic device **804** makes the determination based on the device types broadcasted by the other electronic devices and the list of acceptable device types. In some examples, if electronic device **804** determines that none of the other electronic devices are of an acceptable device type for handling the task, electronic device **804** outputs an error message to the user or prompts the user for additional input (e.g., “would you like the video playback on your iPhone?”).

[0266] Similarly, electronic device **806** (implemented as a television) receives respective sets of one or more values from electronic devices **802** (implemented as a smart watch), **808** (implemented as a speaker dock), and **804** (implemented as a mobile phone). By analyzing the broadcasted sets of values, electronic device **806** determines that it has not broadcasted the highest “energy level” value. But, electronic device **806** further determines that it is of a device type that can handle the task of video playback. In accordance with the determination, electronic device **806** makes an additional determination of whether it is to respond to the audio input despite not having broadcasted the highest “energy level” value. For example, if electronic device **806** determines that none of the electronic devices that have broadcasted higher “energy level” values (e.g. electronic device **804**) are of acceptable device types for handling the task, electronic device **806** responds to the audio input. On the other hand, if electronic device **806** determines that there is at least one electronic device that has broadcasted a higher “energy level” value and is of an acceptable device type for handling the task, electronic device **806** foregoes responding

to the audio input. Forgoing responding to the audio input may include entering an inactive mode in some examples.

[0267] Turning to FIG. 8B, user 800 provides audio input 812, which includes only a spoken trigger (“Hey Siri”) and does not specify a task. In response, each of electronic devices 802-808 generates and broadcasts a set of one or more values, as described. In some examples, each of electronic devices 802-808 broadcasts an “energy level” value to each other. In some examples, each of electronic devices 802-808 calculates a single value aggregating some or all of the exemplary values disclosed herein. In this example, electronic device 804 determines that it has broadcasted the highest “energy level” value, indicating that electronic device 804 is closer to user 800 than the other electronic devices. Accordingly, electronic device 804 responds to the audio input 812, and the rest of the electronic devices forego responding to the audio input.

[0268] In some examples, because audio input 812 (“Hey Siri”) does not specify a task, electronic devices 802-808 may forego broadcasting “device type” values and rely on “energy level” values solely in the arbitration process. In some examples, electronic devices 802-808 forego broadcasting the “device type” values upon determining that the user has not provided another utterance after the spoken trigger (e.g., “Hey Siri”) for a predetermined period of time. It will be appreciated that, after an electronic device determines to respond to audio input 812 (“Hey Siri”), the electronic device can receive an utterance specifying a task and can cause the task to be performed at a second electronic device if the electronic device is not of a type for performing the specified task or is otherwise better suited for performing the specified task.

[0269] Turning to FIG. 8C, user 800 lifts electronic device 802 (implemented as a smart watch) into a raised position and then provides audio input 814 (“Hey Siri”). In some examples, the set of one or more values broadcasted by electronic device 802 includes value(s) indicative of the user’s gesture input at electronic device 802. The value(s) can be a single value indicative of the presence of a recent user input, or can include detailed information regarding the user input, such as a time stamp, an input type, etc. In view of the broadcasted “state” value(s) of electronic device 802, electronic devices 804-808 forego responding to the audio input and electronic device 802 responds to the audio input. In some examples, an electronic device that has broadcasted value(s) indicative of a user input of a particular type (e.g., an active gesture, a physical touch) at the electronic device is to respond to the audio input, regardless of the other values broadcasted. In some examples, when multiple electronic devices have broadcasted values indicative of user inputs of particular types at the respective devices, the electronic device that has broadcasted a value (e.g., time stamp) indicative of the most recent user input is to respond to the audio input, regardless of the other values broadcasted.

[0270] In some examples, if an electronic device broadcasts a set of one or more values and does not receive any set of one or more values from another electronic device, the electronic device responds to the audio input regardless of the values broadcasted.

5. Exemplary Techniques for Intelligent Device Control

[0271] FIGS. 9A-9C illustrate exemplary techniques for intelligent device control in accordance with some embodi-

ment. These figures are also used to illustrate the processes described below, including the processes in FIGS. 11A-E.

[0272] FIG. 9A illustrates an exemplary system and environment 900 for controlling electronic devices using a virtual assistant. A first user device 908 and a second user device 910 may be any of devices 104, 122, 200, 400, 600, and 1330 (FIGS. 1, 2A, 3, 4, 5A, 6A-6B, and 13) in some embodiments. In some examples, each of user devices 908 and 910 may be a television, a set-top box, an audio speaker, a mobile phone, a smart watch, a laptop, a desktop computer, a tablet, or a combination thereof. Virtual assistant server 902 may be implemented using any system described herein, such as system 108 of FIG. 1. Each of media identification storage 906 and device identification storage 904 may reside on virtual assistant server 902, the first user device 908, the second user device 910, or any other device or system in communication with virtual assistant server 902.

[0273] As illustrated in FIG. 9A, in operation, the first user device 908 receives a user input (e.g., “Play The Ramones in my living room”). In some examples, the user input includes user speech and the first user device 908 receives the user input using a microphone of the first electronic device 908. In some examples, the first electronic device 908 converts the user input into a representation of the audio input. The representation of the audio input may be an analog or digital representation. In some examples, the representation is a textual representation and the first electronic device 908 converts the user input into a textual representation using speech-to-text conversion. The user speech can be converted using any known speech-to-text conversation process.

[0274] At 920, the user input, or a representation thereof, is provided (e.g., transmitted) to virtual assistant server 902 by the first electronic device 908 for processing.

[0275] Based on the user input, the virtual assistant server 902 identifies one or more tasks and parameters. For instance, virtual assistant server 902 interprets the textual representation of the user’s input to derive an intent and operationalizes the intent into one or more tasks. In the depicted example, based on the user input “Play The Ramones in my living room”, virtual assistant server 902 identifies the task to be adding one or more media items to a media playback queue at a desired electronic device. The task can be identified in any number of ways. Techniques for identifying tasks can be found, for example, in U.S. Utility application Ser. No. 14/503,105, entitled “Intelligent Assistant for Home Automation”, filed Sep. 30, 2014, and U.S. Provisional Patent Application Ser. No. 62/348,015, entitled “Intelligent Automated Assistant in a Home Environment,” filed Jun. 9, 2016 (Attorney Docket No. 770003000100 (P30331USP1)), the entire disclosures of which are incorporated herein by reference.

[0276] Further, virtual assistant server 902 identifies relevant parameters needed to complete the task(s). In some examples, the parameters include an indication of media item(s) to be played, such as song titles and media file names. In some examples, the parameters include an indication of the type of media item(s) to be played, such as audio, video, and text. In some examples, the parameters include an indication of a location of the desired electronic device, such as named locations and configurations associated with a user (“living room”, “garage”, “upstairs”) or GPS coordinates. In some examples, the parameters include an indication of a type of the desired electronic device, such

as television, speaker dock, and smart watch. In the depicted example, virtual assistant server 902 identifies the parameters to be “Ramones” and “LivingRoom”.

[0277] At 940, virtual assistant server 902 transmits a request for identification of one or more user devices to device identification storage 904. The request includes one or more parameters for identifying the one or more user devices. In some examples, the one or more parameters may specify a named location (e.g., “LivingRoom”), a type of device (e.g., “AppleTV”), and/or one or more attributes of a media item, such as the type of the media item (e.g., “audio” and “video”).

[0278] Device identification storage 904 can store data and/or models for a variety of attributes and relevant information associated with various electronic devices. Attributes that can be stored for an electronic device include, but are not limited to, a unique identifier, a state, a type, a location, media types supported, and ownership. As an example, device identification storage 904 may store information associating a device with a type “AppleTV”, a named location “LivingRoom”, and supported media types “audio” and “video”. As another example, device identification storage 904 may store information associating a device with a type “speaker”, a named location “garage”, and supported media type “audio”.

[0279] Device identification storage 904 may be implemented via hardware, software, or a combination thereof. In some examples, device identification storage 904 may reside on one or more of the above-described electronic devices including the first user device 908 and the second user device 910. Device identification storage 904 may reside on the same electronic device(s) that host any of media identification storage 906 and virtual assistant server 902.

[0280] At device identification storage 904, one or more user devices are identified based on the parameter(s) provided in the request from virtual assistant server 902. In some examples, device identification storage 804 determines whether any of the electronic devices having an entry in the database has attributes matching the parameter(s). In the depicted example, user device 910 is identified by device identification storage 904 based on the request received at 940. Additional details regarding the identification of electronic devices can be found, for example, in U.S. Utility application Ser. No. 14/503,105, entitled “Intelligent Assistant for Home Automation”, filed Sep. 30, 2014, the entire disclosures of which is incorporated herein by reference.

[0281] In response to the request for identification of one or more devices, at 942, virtual assistant server 902 receives identification information of the one or more devices. In the depicted example, the identification information includes a unique identifier corresponding to the identified user device 910 (e.g., “uid=123456”).

[0282] At 960, virtual assistant server 902 transmits a request for identification of one or more media items to media identification storage 906. The request includes one or more parameters allowing for identification of the one or more media items, such as the parameter “Ramones”. Media identification storage 906 includes a database relating (e.g., cataloging) a set of media items by one or more attributes, including title, artist, genre, length, and any other relevant information associated with the media items. By way of example, one or more media items are identified based on the parameter “Ramones” at media identification storage 906.

[0283] In response to the request for identification of one or more media items, at 962, virtual assistant server 902 receives identification information of the one or more media items identified from media identification storage 906. In some examples, the identification information includes a unique identifier corresponding to the identified one or more media items, such as “artist://store/56989”. The unique identifier of the one or more media items can be used to identify, and thereafter access, the one or more identified media items, for instance using one or more application supporting media playback. Media identification storage 906 may reside on any local or remote electronic device(s), such as the same electronic device(s) where virtual assistant server 902 resides. It will be appreciated that while operation at 940 is described as occurring prior to operation at 960, operation corresponding to 960 and 940 may be performed in any chronological order or simultaneously.

[0284] At 980, virtual assistant server 902 provides a unified command to the first user device 908. The command includes the identification information of the one or more media items and the identification information of user device 910. In the depicted example in FIG. 9A, the unified command includes a text string “SetMediaPlayerQueue artist://store/56989 AirplayRouteUid=12456”. In response to receiving the unified command, at 982, user device 908 causes playback of the identified one or more media items (e.g., media pertaining to the Ramones) at the identified user device (e.g., second user device 910, implemented as a television). In some examples, the first user device 908 transmits the relevant information derived from the unified command to the second user device 910.

[0285] With reference to FIG. 9B, in some examples, device identification storage 904 provides (e.g., generates) and/or updates entries in its database of electronic devices by communicating with device tracking storage 912. Device tracking storage 912 provides for the discovery of electronic devices having certain attributes and communicates with device identification storage 904 to update the database of device identification storage 904. Device tracking storage 912 may be implemented via hardware, software, or a combination thereof. Furthermore, device tracking storage 912 may reside on any of the electronic devices described herein, such as first user device 908 or second user device 910. Device tracking storage 912 may reside on the same electronic device(s) that host any of media identification storage 906, virtual assistant server 902, device identification storage 904, or any device or system in communication with device identification storage 904.

[0286] In some examples, device identification storage 904 provides and/or updates entries for electronic devices in its database in response to a user input (e.g., “Play the Ramones on my Apple TV”). As illustrated in FIG. 9B, in response to receiving a request for identification of one or more user devices (940), device identification storage 904 communicates with device tracking storage 912 at 930 to cause discovery of all electronic devices having one or more attributes as specified in or derived from the request. In some examples, discovery of devices is performed responsive to the device identification storage 904 failing to identify a device having parameters matching those provided in the request from virtual assistant server 902. At 932, device tracking storage 912 passes a list of newly discovered electronic devices to device identification storage 904. Based on the discovered electronic devices and the corre-

sponding attributes, device identification storage **904** updates the relevant entries in its database. In some examples, data corresponding to the discovered electronic devices and the corresponding attributes is at least partially cached, either at device tracking storage **912** or device identification storage **904**, for a period of time (e.g., a few minutes).

[0287] In some examples, device identification storage **904** automatically and/or periodically provides and/or updates entries for electronic devices in its database. For example, device tracking storage **912** can facilitate a discovery of electronic devices without a request from virtual assistant server **902** and without any communication from device identification storage **904**.

[0288] In some examples, after sending a request to device identification storage **904** at **940**, virtual assistant server **902** receives identification information that requires input from the user. In one example, user input is required in response to device identification storage **904** having identified multiple user devices corresponding to the parameters of the request. As such, the user is prompted to provide a disambiguation input by, in some instances, selecting a device out of the multiple identified user devices. As another example, user input is required to authenticate devices which are, for instance, discovered in response to the device identification storage **904** failing to identify devices matching parameters of a request.

[0289] As illustrated in FIG. 9C, virtual assistant server **902** receives identification information of multiple identified electronic devices from device identification storage **904** at **942**. At **944**, virtual assistant server **902** provides (e.g., transmits) the identification information to user device **908** for user disambiguation. Accordingly, the user device **908** presents (e.g., displays) a list of the multiple identified electronic devices to the user and prompts the user to provide an input including a selection of a device from the list. In some examples, the device list includes differentiating attributes of the devices, such as color and model.

[0290] An indication of the selected device is transmitted from user device **908** to virtual assistant server **902** at **946**. Based on the received data, virtual assistant server **902** generates a unified command and provides the command to user device **908** at **980**. The command includes identification information of one or more media items and identification information of the device selected by the user. As an example, in FIG. 9C, the command is a text string "SetMediaPlayerQueue artist://store/56989 AirplayRouteUId=12456".

[0291] In some examples, virtual assistant server **902** receives identification information of one or more electronic device that requires user authentication. As discussed above, device identification storage **904** may fail to identify user devices with attributes matching the parameters in a request provided by virtual assistant server **902** and, in response, may identify devices by way of discovery. The virtual assistant server **902** provides (e.g., transmits) identification information of devices identified in this manner to user device **908** for user authentication.

[0292] At user device **908**, the user is prompted to authenticate the candidate device and/or provide relevant information to be associated with the candidate device. For example, the user may be prompted with an output (e.g., audio output) "Is this speaker in your living room?" In response, user device **908** receives an input from the user confirming or

rejecting the candidate device and transmits data corresponding to the user's input to virtual assistant server **902**. If the user confirms the candidate device, virtual assistant server **902** generates a unified command accordingly and provides the command to user device **908**, as described. If the user rejects the candidate device, user device **908** outputs an error message, and, in some instances, communicates the user's rejection of candidate device to virtual assistant server **902**.

[0293] In some examples, the database of device identification storage **904** is updated in response to the user's input at user device **908**. The virtual assistant server **902** can, for instance, receive the user's input and, in response, cause the device identification storage **904** to update the database of electronic devices. By way of example, virtual assistant server **902** receives a disambiguation response at **946** and, at **948**, cause entry of the user-selected device to be created and/or updated at the database of device identification storage **904**. For example, the user-selected device is thereafter associated with an existing configuration (e.g., the "living room" configuration) based on the user input. Additional details regarding the creation and management of configurations can be found, for example, in U.S. Utility application Ser. No. 14/503,105, entitled "Intelligent Assistant for Home Automation", filed Sep. 30, 2014, the entire disclosures of which are incorporated herein by reference.

6. Processes for Intelligent Device Arbitration

[0294] FIGS. 10A-C illustrate process **1000** for operating a digital assistant according to various examples. Process **1000** is performed, for example, using one or more electronic devices (e.g., devices **104**, **106**, **200**, **400**, or **600**) implementing a digital assistant. In some examples, the processes are performed at a client-server system (e.g., system **100**) implementing a digital assistant. The blocks of the processes can be divided up in any manner between the server (e.g., DA server **106**) and the client (e.g., user device **104**). In process **1000**, some blocks are, optionally, combined, the order of some blocks is, optionally, changed, and some blocks are, optionally, omitted. In some examples, only a subset of the features or blocks described below with reference to FIGS. 10A-C is performed.

[0295] As described below, method **1000** provides an efficient way for arbitrating among multiple devices for responding a user input. The method reduces the cognitive burden on a user for managing multiple devices, thereby creating a more efficient human-machine interface. For battery-operated computing devices, enabling a user to reach the completion of tasks more efficiently conserves power and increases the time between battery charges.

[0296] At block **1002**, a first electronic device having a microphone samples an audio input. In some examples, the audio input can be received via a microphone (e.g., microphone **213**) of the electronic device. The received audio input can be processed (e.g., using audio circuitry **210** or processor(s) **220**) and converted into a representative form such as, for example, an electronic signal (digital or analog) or one or more audio files.

[0297] At block **1008**, the first electronic device broadcasts a first set of one or more values based on the sampled audio input. In some examples, a value of the first set of values is based on a signal to noise ratio of speech of the audio input sampled with the first electronic device. In some examples, a value of the first set of values is based on a

sound pressure of the audio input sampled with the first electronic device. In some examples, the first electronic device identifies a confidence value indicative of a likelihood that the audio input was provided by a particular user, and a value of the first set of values is based on the confidence value. In some examples, the first electronic device identifies a state of the first electronic device, and a value of the first set of values is based on the identified state of the first electronic device. In some examples, the state of the first electronic device is identified based on a user input received with the first electronic device. In some examples, at least one value of the first set of one or more values is based on a type of the first electronic device.

[0298] At block 1004, the first electronic device optionally determines whether the audio input comprises a spoken trigger (e.g., “Hey Siri”). At block 1006, in accordance with a determination that the audio input does not comprise the spoken trigger, the first electronic device foregoes broadcasting the first set of one or more values.

[0299] At block 1010, the first electronic device receives a second set of one or more values from a second electronic device. The second set of one or more values is based on the audio input sampled at the second electronic device.

[0300] At block 1012, the first electronic device determines whether the first electronic device is to respond to the audio input based on the first set of one or more values and the second set of one or more values. In some examples, the first electronic device determines whether the first electronic device is to respond to the audio input by determining whether a value of the first set of one or more values is higher than a corresponding value of the second set of one or more values, at block 1042.

[0301] At block 1014, in accordance with a determination that the first electronic device is to respond to the audio input, the first electronic device responds to the audio input. In some examples, the audio input comprises an additional input indicative of a task. At block 1018 (FIG. 10B), the first electronic device additionally determines whether a type of the first electronic device meets a requirement of the task. At block 1020, the first electronic device optionally receives data indicative of the requirement of the task from a server.

[0302] At block 1022, in accordance with a determination that the type of the first electronic device meets the requirement of the task, the first electronic device responds to the audio input. At block 1024, in accordance with a determination that the type of the first electronic device does not meet the requirement of the task, the first electronic device may forego responding to the audio input, or optionally, makes an additional determination of whether the second device is to respond to the audio input. At block 1026, in accordance with a determination that the second device is to respond to the audio input, the first electronic device foregoes responding to the audio input with the first electronic device. At block 1028, in accordance with a determination that the second device is not to respond to the audio input, the first electronic device provides an output indicative of an error.

[0303] At block 1016 (FIG. 10A), in accordance with a determination that the first electronic device is not to respond to the audio input, the first electronic device foregoes responding to the audio input. In some examples, the audio input comprises an additional input indicative of a task. In block 1030 (FIG. 10C), the first electronic device optionally determines whether a type of the first electronic

device meets a requirement of the task. At block 1032, the first electronic device optionally receives data indicative of the requirement of the task from a server.

[0304] At block 1034, in accordance with a determination that the type of the first electronic device meets the requirement of the task, the first electronic device determines whether to respond to the audio input with the first electronic device based on the first set of one or more values, the second set of one or more values, and the requirement of the task. At block 1036, in accordance with a determination to respond to the audio input with the first electronic device, the first electronic device responds to the audio input with the first electronic device. At block 1038, in accordance with a determination to not respond to the audio input with the first electronic device, the first electronic device foregoes responding to the audio input with the first electronic device.

[0305] At block 1040, in accordance with a determination that the type of the first electronic device does not meet the requirement of the task, the first electronic device foregoes responding to the audio input with the first electronic device. In some examples, foregoing responding to the audio input with the first electronic device comprises entering an inactive mode, as illustrated in block 1044.

[0306] In some examples, the first set of one or more values is broadcasted in accordance with a unidirectional broadcast communications protocol, such as BTLE advertising mode.

[0307] In some examples, in accordance with the determination that the first electronic device is to respond to the audio input, the first electronic device provides a visual output, an auditory output, a haptic output, or a combination thereof.

7. Processes for Intelligent Device Control

[0308] FIGS. 11A-E illustrate process 1100 for operating a digital assistant according to various examples. Process 1100 is performed, for example, using one or more electronic devices (e.g., devices 104, 106, 200, 400, 600, 1300, 1330, or 1360) implementing a digital assistant. In some examples, the processes are performed at a client-server system (e.g., system 100) implementing a digital assistant. The blocks of the processes can be divided up in any manner between the server (e.g., DA server 106) and the client (e.g., user device 104). In process 1100, some blocks are, optionally, combined, the order of some blocks is, optionally, changed, and some blocks are, optionally, omitted. In some examples, only a subset of the features or blocks described below with reference to FIGS. 11A-E is performed.

[0309] As described below, method 1100 provides an efficient way for controlling one among multiple devices for responding a user input. The method reduces the cognitive burden on a user for managing multiple devices, thereby creating a more efficient human-machine interface. For battery-operated computing devices, enabling a user to reach the completion of tasks more efficiently conserves power and increases the time between battery charges.

[0310] At block 1102, an electronic device receives data corresponding to an audio input from a first user device.

[0311] At block 1104, the electronic device obtains an identification of a second user device based on the data corresponding to the audio input. In some examples, the electronic device is a first electronic device and communicates with a second electronic device to obtain the identification of the second user device. In some examples, the

second electronic device is the first user device (e.g., implemented using the same hardware).

[0312] Optionally, at block 1106, the first electronic device identifies a task and a parameter based on the data corresponding to the audio input. At block 1108, the first electronic device transmits to a second electronic device a request, wherein the request includes the parameter identifying a user device to perform the task.

[0313] At block 1110, the second electronic device receives the request including the parameter. At block 1112, the second electronic device, in response to receiving the request, discovers one or more user devices. Optionally at block 1113, the second electronic device discovers a user device associated with an indication of a named location of a user. At block 1114, the second electronic device updates a record based on attributes of the one or more discovered user devices. The record stores a plurality of sets of attributes respectively corresponding to a plurality of user devices. At block 1116, the second electronic device caches attributes of the discovered one or more user devices. In some examples, the steps in blocks 1112-1116 are automatically and/or periodically performed without any request from the first electronic device. At block 1118, the second electronic device obtains data from the record. In some examples, a user device of the plurality of user devices is a computer, a television, a set-top box, an audio speaker, or a phone. At block 1120 (FIG. 11B), the second electronic device determines the identification of the second user device based on the parameter, where the second user device is part of the plurality of user devices. In some examples, the second electronic device determines the identification of the second user device by determining whether at least one attribute of the second user device matches the parameter, at block 1121. At block 1122, the second electronic device transmits the identification of the second user device to the first electronic device.

[0314] At block 1124, the first electronic device receives the identification of the second user device from the second electronic device. In some examples, the parameter includes an indication of a named location of a user. In some examples, the parameter includes an indication of a type of a device. In some examples, the parameter includes one or more attributes of the one or more media items.

[0315] Optionally, at block 1126, the first electronic device identifies a task and a parameter based on the data corresponding to the audio input. At block 1132, the first electronic device receives a plurality of identifications corresponding to a plurality of user devices from the second electronic device. In some examples, the obtaining of the plurality of identifications at the second electronic device is similar to the steps described with respect to blocks 1108-1122. At block 1134, the first electronic device transmits the plurality of identifications corresponding to the plurality of user devices to the first user device. At block 1136, the first user device receives the plurality of identifications corresponding to the plurality of user devices. At block 1138, the first user device presents the plurality of user devices as selectable options. At block 1140 (FIG. 11C), the first user device receives an input representing the selection of the particular user device of the plurality of user devices. At block 1142, the first user device transmits the data corresponding to the selection of the particular user device to the

first electronic device. In some examples, each of the plurality of user devices has an attribute matching the parameter.

[0316] At block 1144, the first electronic device receives data corresponding to a selection of a particular user device of the plurality of the user devices from the first user device. At block 1146, the first electronic device obtains the identification of the second user device based on the received data corresponding to the selection of the particular user device of the plurality of the user devices.

[0317] Optionally, at block 1148, the first electronic device identifies a task and a parameter based on the data corresponding to the audio input. At block 1150, the second electronic device receives a request for identification of a user device from the first electronic device, wherein the request includes the parameter. At block 1152, the second electronic device obtains data from a record, wherein the record stores a plurality of sets of attributes respectively corresponding to a plurality of user devices. At block 1154 (FIG. 11D), the second electronic device obtains data corresponding to the candidate user device based on the parameter, wherein the candidate user device is not part of the plurality of user devices stored in the record. In some examples, obtaining the data corresponding to the candidate user device based on the parameter at block 1193 further comprises determining, with the second electronic device, that no device of the plurality of user devices has an attribute matching with the parameter at block 1151, and discovering, with the second electronic device, the candidate user device. The candidate user device has an attribute matching the parameter at block 1153. At block 1156, the second electronic device transmits data corresponding to the candidate user device to the first electronic device.

[0318] At block 1158, the first electronic device receives data corresponding to a candidate user device from a second electronic device. At block 1160, the first electronic device transmits the data corresponding to the candidate user device to the first user device. At block 1162, the first user device receives the data corresponding to the candidate user device. At block 1164, the first user device presents the candidate user device as a confirmable option. At block 1166, the first user device receives an input representing the confirmation of the candidate user device. At block 1168 (FIG. 11E), the first user device transmits data corresponding to the confirmation of the candidate user device to the first electronic device. In some examples, the first user device receives a rejection of the candidate user device at block 1171. In response, the first user device presents an output indicative of an error at block 1173.

[0319] At block 1170, the first electronic device receives data corresponding to a confirmation of the candidate user device from the first user device. At block 1172, the first electronic device, in response to receiving the data corresponding to the confirmation, obtains the identification of the second user device by obtaining an identification of the candidate user device. In some examples, the second electronic device receives the data corresponding to the confirmation of the candidate user device to the first electronic device at block 1171 and, in response to receiving the data corresponding to the confirmation, updates the record at block 1173.

[0320] At block 1174, the electronic device obtains an identification of a media item based on the data corresponding to the audio input. Optionally, at block 1176, the first

electronic device identifies a task and a parameter based on the data corresponding to the audio input. Optionally, at block 1178, the first electronic device transmits a request to a database, wherein the request includes the parameter identifying one or more media items. Optionally, at block 1180, the first electronic device receives the identification of the one or more media items from the database. In some examples, the identification of the media item includes a unique identifier of the one or more media items.

[0321] In accordance with some embodiments, FIG. 12 shows a functional block diagram of an electronic device 1200 configured in accordance with the principles of the various described embodiments, including those described with reference to FIGS. 8A-C and 10A-C. The functional blocks of the device are, optionally, implemented by hardware, software, or a combination of hardware and software to carry out the principles of the various described embodiments. It is understood by persons of skill in the art that the functional blocks described in FIG. 12 are, optionally, combined or separated into sub-blocks to implement the principles of the various described embodiments. Therefore, the description herein optionally supports any possible combination or separation or further definition of the functional blocks described herein.

[0322] As shown in FIG. 12, an electronic device 1200 includes a microphone unit 1202, optionally an output unit 1204, and a processing unit 1208 coupled to the microphone unit 1202 and optionally the output unit 1204. In some embodiments, the processing unit 1208 includes a broadcasting unit 1210, a receiving unit 1212, a determining unit 1214, a responding unit 1216, and optionally, identifying unit 1218 and entering unit 1220.

[0323] In some examples, the processing unit 1208 is configured to: sample (e.g., with the microphone unit 1210) an audio input (e.g., block 1002 of FIG. 10A); broadcast (e.g., with the broadcasting unit 1210) a first set of one or more values based on the sampled audio input (e.g., block 1008 of FIG. 10A); receive (e.g., with the receiving unit 1212) a second set of one or more values from a second electronic device, wherein the second set of one or more values are based on the audio input (e.g., block 1010 of FIG. 10A); determine (e.g., with the determining unit 1214) whether the electronic device is to respond to the audio input based on the first set of one or more values and the second set of one or more values (e.g., block 1012 of FIG. 10A); in accordance with a determination that the electronic device is to respond to the audio input, respond (e.g., with the responding unit 1216) to the audio input (e.g., block 1014 of FIG. 10A); and in accordance with a determination that the electronic device is not to respond to the audio input, forego responding (e.g., with the responding unit 1216) to the audio input (e.g., block 1016 of FIG. 10A).

[0324] In some examples, a value of the first set of values is based on a signal to noise ratio of speech of the audio input sampled with the electronic device 1200.

[0325] In some examples, a value of the first set of values is based on a sound pressure of the audio input sampled with the electronic device 1200.

[0326] In some examples, the processing unit 1208 is further configured to identify (e.g., using the identifying unit 1218) a confidence value indicative of a likelihood that the audio input was provided by a particular user, and a value of the first set of values is based on the confidence value.

[0327] In some examples, the processing unit 1208 is further configured to identify (e.g., using the identifying unit 1218) a state of the electronic device 1200, and a value of the first set of values is based on the identified state of the electronic device 1200.

[0328] In some examples, the state of the electronic device 1200 is identified based on a user input received with the electronic device 1200.

[0329] In some examples, at least one value of the first set of one or more values is based on a type of the electronic device 1200.

[0330] In some examples, sampling the audio input comprises determining (e.g., with determining unit 1214) whether the audio input comprises a spoken trigger (e.g., block 1004 of FIG. 10A). The processing unit 1208 is further configured to: in accordance with a determination that the audio input does not comprise the spoken trigger, forego broadcasting (e.g., with the broadcasting unit 1210) the first set of one or more values.

[0331] In some examples, the audio input comprises an additional input indicative of a task, and responding to the audio input with the electronic device 1200 further comprises: determining (e.g., with the determining unit 1214) whether a type of the electronic device meets a requirement of the task (e.g., block 1018 of FIG. 10B); in accordance with a determination that the type of the electronic device meets the requirement of the task, respond (e.g., with the responding unit 1216) to the audio input (e.g., block 1022 of FIG. 10B); in accordance with a determination that the type of the electronic device does not meet the requirement of the task, forego responding (e.g., with the responding unit 1216) to the audio input (e.g., block 1024 of FIG. 10B).

[0332] In some examples, the processing unit 1208 is further configured to: in accordance with the determination that the type of the electronic device does not meet the requirement, determine (e.g., with the determining unit 1214) whether the second device is to respond to the audio input (e.g., block 1024 of FIG. 10B); in accordance with a determination that the second device is to respond to the audio input, forego responding (e.g., with the responding unit 1216) to the audio input with the electronic device (e.g., block 1028 of FIG. 10B); in accordance with a determination that the second device is not to respond to the audio input, provide (e.g., with the output unit 1204) an output indicative of an error (e.g., block 1026 of FIG. 10B).

[0333] In some examples, the processing unit 1208 is further configured to: receive (e.g., with the receiving unit 1212) data indicative of the requirement of the task from a server (e.g., block 1020 of FIG. 10B).

[0334] In some examples, the audio input comprises an additional input indicative of a task, and foregoing responding to the audio input with the electronic device further comprises: determining (e.g., with determining unit 1214) whether a type of the electronic device meets a requirement of the task (e.g., block 1030 of FIG. 10C); in accordance with a determination that the type of the electronic device meets the requirement of the task, determining (e.g., with determining unit 1214) whether to respond to the audio input with the electronic device based on the first set of one or more values, the second set of one or more values, and the requirement of the task (e.g., block 1034 of FIG. 10C); in accordance with a determination to respond to the audio input with the electronic device, responding (e.g., with responding unit 1216) to the audio input with the electronic

device (e.g., block 1036 of FIG. 10C); and in accordance with a determination to not respond to the audio input with the electronic device, foregoing responding (e.g., with responding unit 1216) to the audio input with the electronic device (e.g., block 1038 of FIG. 10C); and in accordance with a determination that the type of the electronic device does not meet the requirement of the task, foregoing responding (e.g., with responding unit 1216) to the audio input with the electronic device (e.g., block 1040 of FIG. 10C).

[0335] In some examples, the processing unit 1208 is further configured to: receive (e.g., with the receiving unit 1212) data indicative of the requirement of the task from a server (e.g., block 1032 of FIG. 10C).

[0336] In some examples, foregoing responding to the audio input with the electronic device 1200 comprises entering (e.g., with entering unit 1220) an inactive mode (e.g., block 1044 of FIG. 10A).

[0337] In some examples, the first set of one or more values is broadcasted in accordance with a unidirectional broadcast communications protocol.

[0338] In some examples, the processing unit 1208 is further configured to: in accordance with the determination that the electronic device is to respond to the audio input, provide (e.g., with the output unit 1204) a visual output, an auditory output, a haptic output, or a combination thereof.

[0339] In some examples, determining whether the electronic device is to respond to the audio input comprises: determining (e.g., with determining unit 1214) whether a value of the first set of one or more values is higher than a corresponding value of the second set of one or more values (e.g., block 1042 of FIG. 10A).

[0340] The operations described above with respect to FIG. 10A-C are, optionally, implemented by components depicted in FIGS. 1-4, 6A-B, 7A, and 12. For example, sampling operation 1002, determining operation 1004, broadcasting operation 1006, and receiving operation 1010 are optionally implemented by processor(s) 120. It would be clear to a person of ordinary skill in the art how other processes can be implemented based on the components depicted in FIGS. 1-4, 6A-B, 7A, and 12.

[0341] In accordance with some embodiments, FIG. 13 shows a functional block diagram of a system comprising a first electronic device 1300, a first user device 1330, and a second electronic device 1360, at least one of which is configured in accordance with the principles of the various described embodiments, including those described with reference to FIGS. 9A-C and 11A-E. The functional blocks of the device are, optionally, implemented by hardware, software, or a combination of hardware and software to carry out the principles of the various described embodiments. It is understood by persons of skill in the art that the functional blocks described in FIG. 13 are, optionally, combined or separated into sub-blocks to implement the principles of the various described embodiments. Therefore, the description herein optionally supports any possible combination or separation or further definition of the functional blocks described herein.

[0342] As shown in FIG. 13, a first electronic device 1300 includes a first processing unit 1308. In some embodiments, the processing unit 1308 includes a receiving unit 1310, an obtaining unit 1312, a providing unit 1314, and optionally, an identifying unit 1316 and a transmitting unit 1318. A first user device 1330 includes a second processing unit 1338

optionally coupled to an optional output unit. The second processing unit 1338 optionally includes a causing unit 1340, a receiving unit 1342, a determining unit 1344, a responding unit 1346, an identifying unit 1348, and a transmitting unit 1350. A second electronic device 1360 includes a third processing unit 1368. The third processing unit 1368 optionally includes a receiving unit 1370, an obtaining unit 1372, a determining unit 1374, a transmitting unit 1376, a discovering unit 1378, an updating unit 1380, and a caching unit 1382.

[0343] In some examples, the first processing unit 1308 is configured to: receive (e.g., with receiving unit 1330) data corresponding to an audio input from the first user device 1330 (e.g., block 1102 of FIG. 11A); obtain (e.g., with obtaining unit 1312) an identification of a second user device based on the data corresponding to the audio input (e.g., block 1104 of FIG. 11A); obtain (e.g., with obtaining unit 1312) an identification of a media item based on the data corresponding to the audio input (e.g., block 1174 of FIG. 11E); provide (e.g., with providing unit 1314) a command including the identification of the media item and the identification of the second user device to the first user device 1330 (e.g., block 1182 of FIG. 11E).

[0344] In some examples, the second processing unit 1338 is configured to cause playback of the media item at the second user device (e.g., block 1184 of FIG. 11E).

[0345] In some examples, obtaining the identification of the media item comprises: identifying (e.g., with the identifying unit 1316) a task and a parameter based on the data corresponding to the audio input (e.g., block 1176 of FIG. 11E); transmitting (e.g., with the transmitting unit 1318) a request to a database, wherein the request includes the parameter identifying one or more media items (e.g., block 1178 of FIG. 11E); and receiving (e.g., with the receiving unit 1310) the identification of the one or more media items from the database (e.g., block 1180 of FIG. 11E).

[0346] In some examples, the identification of the media item includes a unique identifier of the one or more media items.

[0347] In some examples, obtaining the identification of the second user device comprises: identifying (e.g., with the identifying unit 1316) a task and a parameter based on the data corresponding to the audio input (e.g., block 1106 of FIG. 11A); transmitting (e.g., with the transmitting unit 1318) to the second electronic device 1360 a request, wherein the request includes the parameter identifying a user device to perform the task (e.g., block 1108 of FIG. 11A); and receiving (e.g., with the receiving unit 1310) the identification of the second user device from the second electronic device 1360 (e.g., block 1124 of FIG. 11B).

[0348] In some examples, the parameter includes an indication of a named location of a user.

[0349] In some examples, the parameter includes an indication of a type of a device.

[0350] In some examples, the parameter includes one or more attributes of the one or more media items.

[0351] In some examples, the second electronic device 1360 is the first user device 1330.

[0352] In some examples, the third processing unit 1368 is configured to: receive (e.g., with receiving unit 1370) the request including the parameter (e.g., block 1110 of FIG. 11A); obtain (e.g., with obtaining unit 1372) data from a record, wherein the record stores a plurality of sets of attributes respectively corresponding to a plurality of user

devices (e.g., block **1118** of FIG. **11A**); determine (e.g., with the determining unit **1374**) the identification of the second user device based on the parameter, wherein the second user device is part of the plurality of user devices (e.g., block **1120** of FIG. **11B**); and transmit (e.g., with transmitting unit **1376**) the identification of the second user device to the first electronic device **1300** (e.g., block **1122** of FIG. **11B**).

[0353] In some examples, the determining of the identification of the second user device based on the parameter comprises: determining (e.g., with the determining unit **1374**) whether at least one attribute of the second user device matches the parameter (e.g., block **1121** of FIG. **11B**).

[0354] In some examples, a user device of the plurality of user devices is a computer.

[0355] In some examples, a user device of the plurality of user devices is a television.

[0356] In some examples, a user device of the plurality of user devices is a set-top box.

[0357] In some examples, a user device of the plurality of user devices is an audio speaker.

[0358] In some examples, a user device of the plurality of user devices is a phone.

[0359] In some examples, the third processing unit **1368** is further configured to: discover (e.g., with discovering unit **1378**) one or more user devices and update (e.g., with updating unit **1380**) the record based on attributes of the one or more discovered user devices. These steps are similar to those performed at blocks **1112** and **1114** (FIG. **11A**), but are performed automatically and/or periodically (as opposed to in response to receiving a user request).

[0360] In some examples, the third processing unit **1368** is further configured to: in response to receiving the request including the parameter, discover (e.g., with discovering unit **1378**) one or more user devices (e.g., block **1112** of FIG. **11A**) and update (e.g., with updating unit **1380**) the record based on attributes of the one or more discovered user devices (e.g., block **1114** of FIG. **11A**).

[0361] In some examples, the discovering of the one or more user devices comprises discovering a user device associated with an indication of a named location of a user (e.g., block **1113** of FIG. **11A**).

[0362] In some examples, the third processing unit **1368** is further configured to cache (e.g., with the caching unit **1382**) attributes of the discovered one or more user devices (e.g., block **1118** of FIG. **11A**).

[0363] In some examples, obtaining the identification of the second user device based on the data corresponding to the audio input further comprises: identifying (e.g., with the identifying unit **1316**) a task and a parameter based on the data corresponding to the audio input (e.g., block **1126** of FIG. **11B**); receiving (e.g., with the receiving unit **1310**) a plurality of identifications corresponding to a plurality of user devices from the second electronic device **1360** (e.g., block **1132** of FIG. **11B**); transmitting (e.g., with the transmitting unit **1318**) the plurality of identifications corresponding to the plurality of user devices to the first user device **1330** (e.g., block **1134** of FIG. **11B**); receiving (e.g., with the receiving unit **1310**) data corresponding to a selection of a particular user device of the plurality of the user devices from the first user device **1330** (e.g., block **1144** of FIG. **11C**); obtaining (e.g., with the obtaining unit **1312**) the identification of the second user device based on the

received data corresponding to the selection of the particular user device of the plurality of the user devices (e.g., block **1146** of FIG. **11C**).

[0364] In some examples, the second processing unit **1338** is further configured to: receive (e.g., with receiving unit **1342**) the plurality of identifications corresponding to the plurality of user devices (e.g., block **1136** of FIG. **11B**); present (e.g., with output unit **1334**) the plurality of user devices as selectable options (e.g., block **1138** of FIG. **11B**); receive (e.g., with receiving unit **1342**) an input representing the selection of the particular user device of the plurality of user devices (e.g., block **1140** of FIG. **11C**); transmit (e.g., with transmitting unit **1350**) the data corresponding to the selection of the particular user device to the first electronic device **1300** (e.g., block **1142** of FIG. **11C**).

[0365] In some examples, each of the plurality of user devices has an attribute matching the parameter.

[0366] In some examples, obtaining the identification of the second user device based on the data corresponding to the audio input further comprises: identifying (e.g., with the identifying unit **1316**) a task and a parameter based on the data corresponding to the audio input (e.g., block **1148** of FIG. **11C**); receiving (e.g., with the receiving unit **1310**) data corresponding to a candidate user device from the second electronic device **1360** (e.g., block **1158** of FIG. **11D**); transmitting (e.g., with the transmitting unit **1318**) the data corresponding to the candidate user device to the first user device **1330** (e.g., block **1160** of FIG. **11D**); receiving (e.g., with the receiving unit **1310**) data corresponding to a confirmation of the candidate user device from the first user device **1330** (e.g., block **1170** of FIG. **11E**); and in response to receiving the data corresponding to the confirmation, obtaining (e.g., with obtaining unit **1312**) the identification of the second user device by obtaining an identification of the candidate user device (e.g., block **1172** of FIG. **11E**).

[0367] In some examples, the second processing unit **1338** is further configured to: receive (e.g., with receiving unit **1342**) the data corresponding to the candidate user device (e.g., block **1162** of FIG. **11D**); present (e.g., with output unit **1334**) the candidate user device as a confirmable option (e.g., block **1164** of FIG. **11D**); receive (e.g., with receiving unit **1342**) an input representing the confirmation of the candidate user device (e.g., block **1166** of FIG. **11D**); and transmit (e.g., with transmitting unit **1350**) data corresponding to the confirmation of the candidate user device to the first electronic device **1300** (e.g., block **1168** of FIG. **11E**).

[0368] In some examples, the second processing unit **1338** is further configured to: receive (e.g., with receiving unit **1342**) the data corresponding to the candidate user device; present (e.g., with output unit **1334**) the candidate user device as a confirmable option; receive (e.g., with receiving unit **1342**) an input representing a rejection of the candidate user device (e.g., block **1171** of FIG. **11D**); present (e.g., with output unit **1334**) an output indicative of an error (e.g., block **1173** of FIG. **11D**).

[0369] In some examples, the third processing unit **1368** is further configured to: receive (e.g., with receiving unit **1370**) a request for identification of a user device from the first electronic device **1300**, wherein the request includes the parameter (e.g., block **1150** of FIG. **11C**); obtain (e.g., with obtaining unit **1372**) data from a record, wherein the record stores a plurality of sets of attributes respectively corresponding to a plurality of user devices (e.g., block **1152** of FIG. **11C**); obtain (e.g., with obtaining unit **1372**) data

corresponding to the candidate user device based on the parameter, wherein the candidate user device is not part of the plurality of user devices stored in the record (e.g., block 1154 of FIG. 11D); and transmit (e.g., with transmitting unit 1376) data corresponding to the candidate user device to the first electronic device 1300 (e.g., block 1156 of FIG. 11D).

[0370] In some examples, obtaining the data corresponding to the candidate user device based on the parameter further comprises: determining (e.g., with determining unit 1374) that no device of the plurality of user devices has an attribute matching with the parameter (e.g., block 1151 in FIG. 11C); and discovering (e.g., with the discovering unit 1378) the candidate user device, wherein the candidate user device has an attribute matching the parameter (e.g., block 1153 in FIG. 11C).

[0371] In some examples, the third processing unit 1368 is further configured to: receive (e.g., with receiving unit 1370) the data corresponding to the confirmation of the candidate user device to the first electronic device 1300 (e.g., block 1171 of FIG. 11E); in response to receiving the data corresponding to the confirmation, update (e.g., with updating unit 1380) the record (e.g., block 1173 of FIG. 11E).

[0372] The operations described above with respect to FIGS. 11A-E are, optionally, implemented by components depicted in FIGS. 1-4, 6A-B, 7A, and 13. For example, receiving operation 1102, obtaining operation 1104, and providing operation 1182 are optionally implemented by processor(s) 120. It would be clear to a person of ordinary skill in the art how other processes can be implemented based on the components depicted in FIGS. 1-4, 6A-B, 7A, and 13.

[0373] In accordance with some implementations, a computer-readable storage medium (e.g., a non-transitory computer readable storage medium) is provided, the computer-readable storage medium storing one or more programs for execution by one or more processors of an electronic device, the one or more programs including instructions for performing any of the methods or processes described herein.

[0374] In accordance with some implementations, an electronic device (e.g., a portable electronic device) is provided that comprises means for performing any of the methods or processes described herein.

[0375] In accordance with some implementations, an electronic device (e.g., a portable electronic device) is provided that comprises a processing unit configured to perform any of the methods or processes described herein.

[0376] In accordance with some implementations, an electronic device (e.g., a portable electronic device) is provided that comprises one or more processors and memory storing one or more programs for execution by the one or more processors, the one or more programs including instructions for performing any of the methods or processes described herein.

[0377] The foregoing description, for purpose of explanation, has been described with reference to specific embodiments. However, the illustrative discussions above are not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many modifications and variations are possible in view of the above teachings. The embodiments were chosen and described in order to best explain the principles of the techniques and their practical applications. Others skilled in the art are thereby enabled to best utilize the techniques and various embodiments with various modifications as are suited to the particular use contemplated.

[0378] Although the disclosure and examples have been fully described with reference to the accompanying drawings, it is to be noted that various changes and modifications will become apparent to those skilled in the art. Such changes and modifications are to be understood as being included within the scope of the disclosure and examples as defined by the claims.

What is claimed is:

1. An electronic device comprising:

a microphone;

one or more processors;

a memory; and

one or more programs, wherein the one or more programs are stored in the memory and configured to be executed by the one or more processors, the one or more programs including instructions for:

sampling, with the microphone at the electronic device, an audio input;

broadcasting a first set of one or more values based on the sampled audio input;

receiving a second set of one or more values from a second electronic device, wherein the second set of one or more values are based on the audio input;

determining whether the electronic device is to respond to the audio input based on the first set of one or more values and the second set of one or more values;

in accordance with a determination that the electronic device is to respond to the audio input, responding to the audio input; and

in accordance with a determination that the electronic device is not to respond to the audio input, foregoing responding to the audio input.

2. The electronic device of claim 1, wherein a value of the first set of values is based on a signal to noise ratio of speech of the audio input sampled with the first electronic device.

3. The electronic device of claim 1, wherein a value of the first set of values is based on a sound pressure of the audio input sampled with the first electronic device.

4. The electronic device of claim 1, wherein the one or more programs further include instructions for:

identifying, with the first electronic device, a confidence value indicative of a likelihood that the audio input was provided by a particular user, wherein a value of the first set of values is based on the confidence value.

5. The electronic device of claim 1, wherein the one or more programs further include instructions for:

identifying, with the first electronic device, a state of the first electronic device, wherein a value of the first set of values is based on the identified state of the first electronic device.

6. The electronic device of claim 5, wherein the state of the first electronic device is identified based on a user input received with the first electronic device.

7. The electronic device of claim 1, wherein at least one value of the first set of one or more values is based on a type of the first electronic device.

8. The electronic device of claim 1, wherein sampling the audio input comprises determining, with the first electronic device, whether the audio input comprises a spoken trigger and wherein the one or more programs further include instructions for:

in accordance with a determination that the audio input does not comprise the spoken trigger, foregoing broadcasting, with the first electronic device, the first set of one or more values.

9. The electronic device of claim 1, wherein the audio input comprises an additional input indicative of a task, and wherein responding to the audio input with the first electronic device further comprises:

determining, with the first electronic device, whether a type of the first electronic device meets a requirement of the task;

in accordance with a determination that the type of the first electronic device meets the requirement of the task, responding to the audio input;

in accordance with a determination that the type of the first electronic device does not meet the requirement of the task, foregoing responding to the audio input.

10. The electronic device of claim 9, wherein the one or more programs further include instructions for:

in accordance with the determination that the type of the first electronic device does not meet the requirement, determining, with the first electronic device, whether the second device is to respond to the audio input,

in accordance with a determination that the second device is to respond to the audio input, foregoing responding to the audio input with the first electronic device;

in accordance with a determination that the second device is not to respond to the audio input, providing, with the first electronic device, an output indicative of an error.

11. The electronic device of claim 9, wherein the one or more programs further include instructions for:

receiving, with the first electronic device, data indicative of the requirement of the task from a server.

12. The electronic device of claim 1, wherein the audio input comprises an additional input indicative of a task, and wherein foregoing responding to the audio input with the first electronic device further comprises:

determining, with the first electronic device, whether a type of the first electronic device meets a requirement of the task;

in accordance with a determination that the type of the first electronic device meets the requirement of the task:

determining whether to respond to the audio input with the first electronic device based on the first set of one or more values, the second set of one or more values, and the requirement of the task;

in accordance with a determination to respond to the audio input with the first electronic device, responding to the audio input with the first electronic device; and

in accordance with a determination to not respond to the audio input with the first electronic device, foregoing responding to the audio input with the first electronic device; and

in accordance with a determination that the type of the first electronic device does not meet the requirement of the task, foregoing responding to the audio input with the first electronic device.

13. The electronic device of claim 12, wherein the one or more programs further include instructions for:

receiving, with the first electronic device, data indicative of the requirement of the task from a server.

14. The electronic device of claim 1, wherein foregoing responding to the audio input with the first electronic device comprises entering, with the first electronic device, an inactive mode.

15. The electronic device of claim 1, wherein the first set of one or more values is broadcasted in accordance with a unidirectional broadcast communications protocol.

16. The electronic device of claim 1, wherein the one or more programs further include instructions for:

in accordance with the determination that the first electronic device is to respond to the audio input, providing, with the first electronic device, a visual output, an auditory output, a haptic output, or a combination thereof.

17. The electronic device of claim 1, wherein determining whether the first electronic device is to respond to the audio input comprises:

determining, with the first electronic device, whether a value of the first set of one or more values is higher than a corresponding value of the second set of one or more values.

18. A non-transitory computer readable storage medium storing one or more programs, the one or more programs comprising instructions, which when executed by one or more processors of an electronic device with a microphone, cause the electronic device to:

sample, with the microphone at the electronic device, an audio input;

broadcast a first set of one or more values based on the sampled audio input;

receive a second set of one or more values from a second electronic device, wherein the second set of one or more values are based on the audio input; and

determine whether the electronic device is to respond to the audio input based on the first set of one or more values and the second set of one or more values;

in accordance with a determination that the electronic device is to respond to the audio input, respond to the audio input;

in accordance with a determination that the electronic device is not to respond to the audio input, foregoing responding to the audio input.

19. A method comprising:

at a first electronic device having a microphone:

sampling, with the microphone at the first electronic device, an audio input;

broadcasting, with the first electronic device, a first set of one or more values based on the sampled audio input;

receiving, with the first electronic device, a second set of one or more values from a second electronic device, wherein the second set of one or more values are based on the audio input;

determining, with the first electronic device, whether the first electronic device is to respond to the audio input based on the first set of one or more values and the second set of one or more values;

in accordance with a determination that the first electronic device is to respond to the audio input, responding to the audio input;

in accordance with a determination that the first
electronic device is not to respond to the audio
input, foregoing responding to the audio input.

* * * * *