October 02, 2017

ASSISTANT-RELATED TASKS USING CONTEXT FROM MULTIPLE DEVICES

Google Inc.

ASSISTANT-RELATED TASKS USING CONTEXT FROM MULTIPLE DEVICES

ABSTRACT

An interactive assistant computing system, referred to herein as “an interactive assistant,” “a virtual assistant,” “a computational assistant,” or simply an “assistant,” is described that, in addition to performing general task requests, will request information from a respective assistant on one or more other computing devices to provide additional information and/or tasks to the user of the interactive assistant. The interactive assistant in the main device may receive an indication of a request from a user to perform a particular task. The interactive assistant, in addition to performing the particular task, may send requests to a respective assistant on one or more other computing devices for context related to the particular task. The respective assistants may provide the context to the interactive assistant, which may utilize that context to provide additional information to the user. This way, the interactive assistant can provide additional information related to the requested task that would previously only be available if the user was interacting with the other computing devices instead of the device executing the interactive assistant.

DESCRIPTION

Assistants execute on counter-top devices, mobile phones, automobiles, and many other type of computing devices. Assistants output useful information, responds to users’ needs, or otherwise performs certain operations to help users complete real-world and/or virtual tasks. Some such tasks include altering the settings of computing devices that executes the assistant, or using one computing device executing the assistant to alter the settings of a second computing device, such as an automobile or other assistant-enabled computing device, or retrieving

information from a database or a network and provide the information to the user. While performing these tasks may be beneficial to the user, the assistant is generally limited to information present on the device executing the assistant, on a network accessible by the assistant, or the device that the assistant is controlling.

In many cases, a user may wish to utilize the assistant to perform a task (e.g., alter the settings of a device or retrieve information for the user). In a simple case, when the task is to retrieve information for the user, the user may initiate a search (e.g., via an audible or a manual query) for the requested information. The assistant may perform a search in a database, locally and/or via a remote server, to provide the user (e.g., audibly, visually) with search results that correspond to the search. However, this search is typically limited to databases immediately accessible to the assistant (e.g., the local database of the computing device executing the assistant or a database stored on a server on a network). For instance, if the user is wishing to determine how long it will take the user to drive to work, the assistant will generally access a service on a network that uses real-time traffic information to determine optimal directions from the user's current location to the user's place of business. However, additional factors may influence the amount of time that the trip may take the user, but these factors may only be available on other computing devices on the same network but not generally accessible by means other than direct interaction. For instance, the user's vehicle may not have enough gasoline to make the entire trip, and the user may have to stop at a gas station along the route to buy additional gasoline. Without this additional information, the user may not allocate adequate time to the trip.

The computing system shown in Figure 1 below enables a primary computing device to perform a task initiated by a user with information polled from one or more secondary

computing devices in addition to the information available locally or via a service on a network. In a simple case, the primary computing device receives an indication of a task that the user would like the primary computing device to perform. The primary computing device may determine that information available at the secondary computing device may store additional information that could assist the primary computing device in performing the task accurately and completely. The primary computing device may send a request to the secondary computing device over the network for the information, and the secondary computing device may provide the information to the primary computing device. The primary computing device may then incorporate the received information into the product produced as a result of the primary computing device performing the requested task.

In this way, the assistant may perform a task requested by the user without requiring the user to directly access multiple computing devices and manually combine the information without the assistance of the computing devices. By automatically retrieving relevant information from one or more additional computing devices, the assistant reduces the time between the user requesting the performance of the task and the user having a complete answer to the query. Further, the assistant reduces the number of interactions the user must perform in order to derive a complete answer to a single interaction with the primary computing device rather than an interaction with each computing device that has information relevant to the task at hand. The techniques described herein have many applications and use cases. Although an assistant in a mobile phone is described throughout the disclosure, an assistant executing in any computing environment could benefit from use of the techniques described herein.

Throughout the disclosure, examples are described where devices and systems analyze information (e.g., locations, movements, calendars, communications, settings, etc.) associated

with users of the devices and systems only if the devices and systems receive explicit permission from the users to analyze the information. For example, in situations discussed below in which multiple computing devices and/or information server system collects and aggregates contextual information regarding various settings, interactions, events, and/or destinations associated with a user, the user may be provided with an opportunity to control whether the devices and/or systems can collect and make use of the information, and to dictate how the devices and systems present information to the user. Additionally, certain data may be treated in one or more ways before the devices and systems store or use the information so that any personally-identifiable information is removed before storage or use. As such, the individual users maintain control over how information is collected about the user and how the information is used by the device and/or system.

Consider the example system shown in Figure 1 which is configured to provide an assistant in accordance with the techniques described herein. The system includes two computing devices communicating across a network.

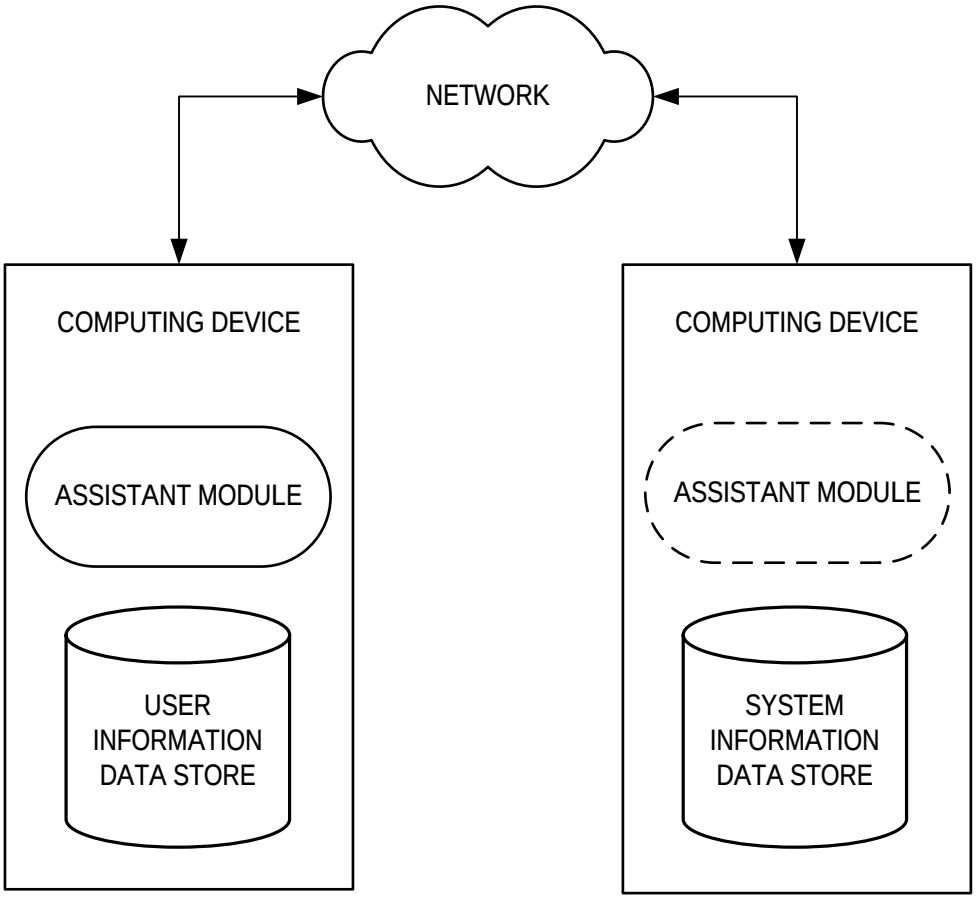

Figure 1

The network represents a combination of any one or more public or private communication networks, for instance, television broadcast networks, short-wavelength wireless networks, cable or satellite networks, cellular networks, Wi-Fi networks, broadband networks, and/or other type of network for transmitting data (e.g., telecommunications and/or media data) between various computing devices, systems, and other communications and media equipment.

The primary computing device (shown with the user information data store and the assistant module) may represent any type of computing device, server, or other system that is configured to execute an assistant and communicate on a network. Examples of the primary

computing device include cloud computing environments, mobile phones, tablet computers, wearable computing devices, countertop computing devices, home automation computing devices, laptop computers, desktop computers, televisions, stereos, automobiles, and all other type of mobile and non-mobile computing device that are configured to execute an assistant.

Similarly, the secondary computing device (shown with the system information data store) may represent any type of computing device, server, cloud computing system, or other system that is configured to communicate on a network with an assistant-enabled device. The secondary computing device may not be assistant-enabled itself, although, in some instances, the secondary computing device may also execute an assistant. In some examples, the secondary computing device may be shared assets of multiple users. For example, the secondary computing device may be a server that stores information regarding the user of the first computing device as well as information regarding other users in the user information data store. The secondary computing device may alternatively be an assistant device closely coupled to another environment, such as a home appliance, automobile, local area network monitoring system, and so on. In such cases, the secondary computing device may have specific information regarding both a user (or multiple users) and the environment itself (referred to as system information), such as temperature or contents within a refrigerator, state of conditions for a vehicle (amount of fuel or charge, tire pressure, etc.), network operational conditions, etc. The secondary computing device may be configured to send the information regarding the user and the user's activities to the first computing device for analysis by the assistant module. The secondary computing device may store system information in the system information data store. The secondary computing device may be configured to send such system information, including

that which is relevant to a task being performed by the primary computing device to the primary computing device when the secondary computing device receives a request for such information.

The assistant module (also referred to as "the assistant") may execute at the primary computing device to provide assistant services to users of the assistant-enabled computing device. Examples of assistant services include altering the settings of computing devices that executes the assistant, or using one computing device executing the assistant to alter the settings of a second computing device, such as an automobile or other assistant-enabled computing device, or retrieving information from a database or a network and provide the information to the user. The user may initiate such tasks by, e.g., providing manual, audible, tactile, or other form of user input.

As a user interacts with the assistant and settings within the primary computing device, the assistant may obtain personal information about the user. Examples of personal information include: habits, preferences for settings within the automobile, location histories, contacts, communications, interests, and other types of user information. After receiving explicit permission from the user, the assistant may store the personal information at the user information data store and, in the course of providing assistant services, make use of the personal information stored at the user information data store.

In these and other examples, while interacting with the user, the assistant may collect, cluster and/or index data for storage in one or more data stores that store interaction history information (e.g., the user information data store). In some cases, the assistant may store the interaction history locally on the assistant-enabled computing device, while in other cases, the assistant may store at least part of the history on one or more external computing systems which

are coupled to the interactive assistant via one or more networks (e.g., one or more wired and/or wireless networks).

The assistant-enabled computing device and the assistant treat the information stored at the user information data store such that the information is protected, encrypted, or otherwise not susceptible to hacking or unauthorized use. The information stored at the user information data store may be stored locally at the assistant-enabled computing device and/or remotely (e.g., in a cloud computing environment accessible via the network). The user information data store may also store information for multiple different users such that the assistant-enabled computing device may be shared amongst multiple users.

As described above, when the primary computing device receives an indication to have the assistant perform a particular task, if the secondary computing device contains information in the system information data store that is relevant to the task being performed by the assistant, then the assistant may query the secondary computing device for such information prior to performing, while performing, or after performing the task. Optionally, the assistant may query the secondary computing device to initiate a measurement or obtain data that is then sent, other than through the data store, to the primary computing device (or to another device, as appropriate to the request from the assistant). For instance, the user may request that the assistant provide the user with an optimal route to get to a particular destination and a time that it will take to reach the particular destination. After determining this information, the assistant may query an assistant in a secondary computing device (i.e., the automobile that the user would use to drive to the particular destination) to determine if the automobile will have enough gasoline to reach the destination. If the automobile has an adequate amount of gasoline, then the primary computing device may output the optimal route and the determined time in an unaltered state. If, however,

the automobile does not have an adequate amount of gasoline, then the primary computing device may include an intermediate stop at a gas station, and the primary computing device may add extra time to the trip duration to account for the intermediate stop.

In another instance, the secondary computing device may be an oven. The primary computing device may receive a request for the assistant to preheat the oven such that it reaches a temperature of 350 degrees Fahrenheit by a particular time. Upon receiving this request, the primary computing device may send a request to the oven for intermittent updates regarding the internal temperature of the oven. The primary computing device may use this information to determine an estimated time at which the primary computing device must send the order to the oven to begin preheating. For example, if the oven is consistently at room temperature, then the primary computing device may determine a first time at which the primary computing device sends the order to the oven to begin the preheating process. Conversely, if the user's spouse recently used the oven and the oven is warmer than room temperature, then the primary computing device may send the order to the oven to begin the preheating process closer to the particular time the oven should be preheated by.

These examples are just two examples of potential tasks that the primary computing device may perform more efficiently and more completely with the use of the additional information provided by a secondary computing device. These techniques may be extended to any task that may be performed by the assistant-enabled primary computing device where having information stored on a secondary computing device and not generally available to the primary computing device may be relevant to the task being performed. Numerous computing devices, including household appliances (e.g., refrigerators, washing machines, televisions, cooking appliances, etc.), vehicles, entertainment centers, tablets, thermostats, and any other device that

may store information relevant to tasks being performed by an assistant on a different computing device may be referenced using the techniques described herein.

In this way, the assistant may perform a task requested by the user without requiring the user to directly access multiple computing devices and manually combine the information without the assistance of the computing devices. By automatically retrieving relevant information from one or more additional computing devices, the assistant reduces the time between the user requesting the performance of the task and the user having a complete answer to the query. Further, the assistant reduces the number of interactions the user must perform in order to derive a complete answer to a single interaction with the primary computing device rather than an interaction with each computing device that has information relevant to the task at hand. The techniques described herein have many applications and use cases. Although an assistant in a mobile phone is described throughout the disclosure, an assistant executing in any computing environment could benefit from use of the techniques described herein.