



- (51) International Patent Classification:
G06F 9/50 (2006.01)
- (21) International Application Number:
PCT/US2017/017090
- (22) International Filing Date:
9 February 2017 (09.02.2017)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
15/048,397 19 February 2016 (19.02.2016) US
- (71) Applicant: MICROSOFT TECHNOLOGY LICENSING, LLC [US/US]; Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US).
- (72) Inventors: KISHAN, Arun U.; MICROSOFT TECHNOLOGY LICENSING, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). WILSON, Emily Nicole; MICROSOFT TECHNOLOGY LICENSING, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). SOLIMAN, Mohammed Nashaat; MICROSOFT TECHNOLOGY LICENSING, LLC, Attn: Patent Group Docketing (Bldg.

8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). MAISURIA, Paresh; MICROSOFT TECHNOLOGY LICENSING, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). WEINBERG, Shira; MICROSOFT TECHNOLOGY LICENSING, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). VIRDI, Gurpreet; MICROSOFT TECHNOLOGY LICENSING, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). BROWN, Jared; MICROSOFT TECHNOLOGY LICENSING, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US).

(74) Agents: MINHAS, Sandip et al.; Microsoft Corporation, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG,

[Continued on next page]

(54) Title: USER PRESENCE PREDICTION DRIVEN DEVICE MANAGEMENT

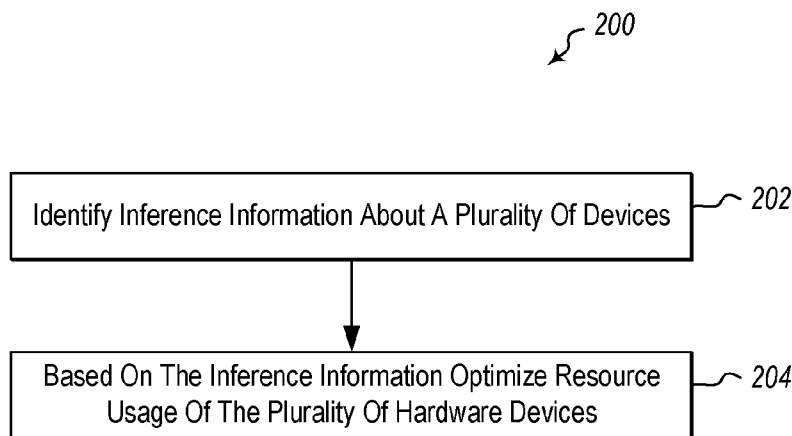


Figure 2

(57) Abstract: Pooling computing resources based on inferences about a plurality of hardware devices. The method includes identifying inference information about the plurality of devices. The method further includes based on the inference information optimizing resource usage of the plurality of hardware devices.



NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE,

Declarations under Rule 4.17:

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

Published:

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

USER PRESENCE PREDICTION DRIVEN DEVICE MANAGEMENT

BACKGROUND

Background and Relevant Art

5 [0001] Computers and computing systems have affected nearly every aspect of modern living. Computers are generally involved in work, recreation, healthcare, transportation, entertainment, household management, etc.

[0002] In the ordinary course of business and leisure, computer users often use multiple computing devices. For example, a computer user may use a cellular smart phone
10 at various times. Contemporaneously, a user may use a laptop or desktop computer. Alternatively or additionally, a user may use a gaming system having large amounts of computing power. Alternatively or additionally, a user may use smart televisions, or other consumer electronic devices, which all have computing power. Each of these devices may also have various peripherals, such as cameras, microphones, screens, speakers, etc.

15 [0003] The subject matter claimed herein is not limited to embodiments that solve any disadvantages or that operate only in environments such as those described above. Rather, this background is only provided to illustrate one exemplary technology area where some embodiments described herein may be practiced.

BRIEF SUMMARY

20 [0004] One embodiment illustrated herein includes a method that may be practiced in a computing environment. The method includes acts for pooling computing resources based on inferences about a plurality of hardware devices. The method includes identifying inference information about the plurality of devices. The method further includes based on the inference information optimizing resource usage of the plurality of
25 hardware devices.

[0005] This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter.

30 [0006] Additional features and advantages will be set forth in the description which follows, and in part will be obvious from the description, or may be learned by the practice of the teachings herein. Features and advantages of the invention may be realized and obtained by means of the instruments and combinations particularly pointed out in the appended claims. Features of the present invention will become more fully apparent from

the following description and appended claims, or may be learned by the practice of the invention as set forth hereinafter.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] In order to describe the manner in which the above-recited and other advantages and features can be obtained, a more particular description of the subject matter briefly described above will be rendered by reference to specific embodiments which are illustrated in the appended drawings. Understanding that these drawings depict only typical embodiments and are not therefore to be considered to be limiting in scope, embodiments will be described and explained with additional specificity and detail through the use of the accompanying drawings in which:

[0008] Figure 1 illustrates a system for pooling resources based on inference information; and

[0009] Figure 2 illustrates a method of pooling resources based on inference information.

DETAILED DESCRIPTION

[0010] Embodiments described herein can identify inference information about a plurality of devices used by a user. From this inference information, embodiments can optimize resource usage of the plurality of hardware devices. For example, by predicting the user's presence across their array of devices (PCs, tablets, phones, videogame consoles, televisions, projectors, etc.), agents running on each of the devices can work together to efficiently pool resources which provide a better experience than any of the devices can offer alone and can save power and energy at the same time.

[0011] For example, consider a case where a user using a phone requests that the phone begin performing some processing activity that will use a large amount of computing resources on the phone. However, the phone knows, using inference information that the user will shortly likely be in close proximity to a desktop computer system. The phone could cause the processing to be delayed as a result of inferring that the user will shortly be proximate the desktop computer system. Once the user is in proximity of the desktop computer system, the processing could be offloaded to the desktop system. This could make the phone appear more efficient and save battery power at the phone as the processing could be offloaded to a desktop computing system that does not rely on battery power. Thus, a phone device with such functionality could be more power efficient than other phone devices.

[0012] Alternatively, consider a case where a user using a phone requests that the phone begin performing some processing activity that will use a large amount of computing resources on the phone. The phone is proximate a desktop computing system, and knows, using inference information that the user will likely stay in close proximity to the desktop computer system for some period of time. The phone could cause the processing to be offloaded to the desktop system, knowing that the desktop computing system will likely be able to complete the processing before the user leaves the proximity of the desktop computing system. This could make the phone appear more efficient and save battery power at the phone as the processing could be offloaded to a desktop computing system that does not rely on battery power.

[0013] Note that embodiments could take into account computing tasks to be performed, relative capabilities of devices, inference about device presence, etc., to determine what device perform computing tasks.

Design

[0014] Some embodiments illustrated herein include three major components:

[0015] (1) user presence inference across devices

[0016] (2) agent modelled device capabilities and resources

[0017] (3) real time understanding of device utilization

[0018] These are now each discussed in more detail below.

User Presence Inference

[0019] Figure 1 illustrates an environment where embodiments may be practiced. In particular, Figure 1 illustrates a personalization server 102. The personalization server 102 may be able to collect data for and about users of devices. For example, Figure 1 also illustrates a user 104. The user 104 may use several different devices over the ordinary course of activities. For example, the user may use a smart cellular telephone device 106-1, a desktop computer device 106-2, a laptop device 106-3, a tablet device 106-4, a gaming device 106-5, a smart television device 106-6 or projector, etc. Each of these devices 106 may have an agent 118-1 through 118-6 that collects data about the user's use of the device 106 respectively, and provides such data to the personalization server 102. Alternatively, the agents may act in a peer to peer fashion such that the personalization server 102 could be eliminated and all such functionality that might be provided by the personalization server would instead be provided by the network of agents. Inference information could be calculated locally at each of the peers, centrally at a managing peer, or in other fashions

[0020] Collected data may include usage pattern information such as information about when a device is used, where a device is used, how heavily a device is used, if the device is used with other devices, the type of activity that the user is engaging in (e.g., navigating towards a destination, starting a lengthy movie, etc.) what actions users take in certain activity contexts (e.g., raising or lowering volumes, silencing devices, activating certain modes, etc.) and any combinations of these to create usage patterns of information, etc. An example of the functionality provided by a personalization solution, such as one that includes the personalization server 102 and/or the agents (referred to generally as 118) is Cortana available from Microsoft Corporation of Redmond, Washington.

[0021] A personalization solution will leverage data from devices 106 to which the user 104 has signed in or otherwise self-identified. By measuring usage of a given devices 106, a model is developed that predicts the user's usage across their devices 106 for any given time, which may include consideration of factors such as day, date, month, year, etc. In particular, agents 118 can measure usage and provide such information to the personalization server 104 or amongst the agents 118.

[0022] Additional contextual user information further improves the model. For example additional context may include the location at which the device 106 is typically used at any given time. Local information is comprised of the user's last input, screen on/off state, remotely connected sessions, etc. Contextual data may further include the user's calendar, their predicted commute period, inactivity based on the usage of a different device and its location (e.g., one's phone) relative to a target device, etc.

[0023] Inferences can include a prediction accuracy or probability factor allowing for deviations in usage patterns. For example, a personalization solution can learn that the user 104 usually uses his work laptop device 106-3 on Monday starting at 8:00am, but starts using it at 11:00am on Tuesday and Wednesday. This could also extend to special events such as different usage patterns during holidays and vacations. However, a given user device usage pattern will likely not be consistent 100% of the time, and thus, a probability factor may be applied to inference information providing an indication of the likelihood that a pattern will occur.

[0024] The inference can also adapt to recognize multiple patterns and adjust the inference at runtime. For example, assume the user 104 has two common usage patterns. In one case, the user 104 is present at their work laptop device 106-3 from 8am to 4pm and then continues working for an hour at 10pm. In another common case, the user is present at their work laptop device from 9am to 7pm with no further usage for that day. Thus, in

some embodiments illustrated in this example, depending how the user starts their day, the inference will adapt dynamically to provide the most likely usage pattern based on the most recent user activity.

Agent Modelled Device Capabilities and Resources

5 [0025] In some embodiments, an agent 118 for a device 106 develops a model 120-1 through 120-6 of its device's capabilities. Both hardware and software (available applications and services) are included in the models 120-1 through 120-6. The models 120-1 through 120-6 are shared between the agents 118-1 through 118-6.

[0026] Simple sensors such as compasses, radios (Cell/LTE, GPS, Wi-Fi, etc.),
10 memory, CPU, hard disk space, speakers, microphones, screens, projectors, etc., can be included in the model.

[0027] Complicated sensors such as infrared cameras, 3D cameras, microphone arrays, holographic screens, 3D printers, etc., can also be included in the model.

[0028] Applications running on the device, accounts, network resources, etc., can also
15 be included in the model.

Real-Time Understanding of Device Utilization

[0029] Agents 118 running on each device 106 monitor the utilization of the device's resources. For example, CPU load, memory utilization, antennae / radio use, SIM presence as well as plan usage (voice, SMS, data), use of specialized hardware (apps running full
20 screen, speakers playing music, etc.), location of the device, user-presence/activity on the device, etc. are all monitored in real-time by agents 118.

[0030] The real-time models 120 are shared between running agents 118 and updated as state changes. In some embodiments, if a device 106 is hibernating, the devices are not woken up to update the model or to receive model updates from agents on other devices
25 rather, when the devices become active, the models can be updated on those devices. However, this may be implemented in other ways as well. For example, in some embodiments, if the models have changed beyond some threshold amount or in some significant way (as defined by a user, administrator, hard coding, or otherwise), then devices may be woken up to receive new model information. Alternatively or additionally,
30 decisions about whether or not to wake up a device may be based on other factors such as whether or not the device is connected to external power, battery charge state of a device, probability that a device is about to be used or may be needed for some computing task, etc.

System

[0031] Based on usage patterns predicted by inference, the capabilities of each device 106 and the real-time utilization signals, the network of agents 118 can make optimal use of the devices resources.

5 [0032] Various examples are now illustrated.

[0033] One example illustrated herein includes functionality for using inference functionality to enable listening functionality using resources from multiple devices 106. In this example, multiple devices' microphones can be pooled to form a broad, ambient microphone array when in close physical proximity. By predicting when those devices are typically near each other, ambient listening applications can adapt to take advantage of the additional hardware the agent makes available. If a microphone becomes used and is no longer available, the model can allow another device to take its place, resulting in
10 continual uptime via seamless hand-off.

[0034] Another example illustrated herein includes functionality for using inference
15 functionality to enable watching functionality using resources from multiple devices 106. For example, some applications such as Windows Hello (available from Microsoft Corporation of Redmond, Washington), XBOX Live sign-in (available from Microsoft Corporation of Redmond, Washington) and potential future usage applications may implement certain levels of privacy based on presence or absence of people within
20 viewing/hearing distance. This may be performed using a system of cameras to identify the user and the presence of other potential users in a given space. By predicting when the user is typically using a given device and which other camera-equipped devices are in the vicinity of that device, a pool of cameras can enable scenarios that are otherwise not possible with just one device. For example, a Windows Hello capable PC could wake up
25 around the time the user typically arrives at a location with an ambient Cortana device that has no camera at all and enable seamless logon to the ambient Cortana device.

[0035] Another example illustrated herein includes functionality for using inference functionality to enable computing functionality using resources from multiple devices 106. For example, should one device's hardware sensors be available but its CPU is occupied
30 with other tasks, the raw data can be streamed to a device with spare cycles to perform the computation. By understanding when the user is likely to consume resources, a secure tunnel between the devices can be pre-established ensuring little or no latency due to secure handshake or tunnel setup.

[0036] In another example, embodiments may be able to perform adaptive power saving operations based on inference information. For example, certain fitness tracker devices may be GPS devices embedded within them. A user may be out exercising with their fitness tracker device. Inference information may predict that a user is about to enter a wooded area where GPS functionality will not function (or that will not function well) and that the user is likely to remain in the wooded area for a given period of time. Using this information, embodiments could disable the GPS for some time based on the given period of time. This would result in energy savings for the fitness tracker device.

[0037] Similarly, inference information may be used to predict that a user is about to embark on a long run with both their fitness tracker device and their phone device. Due to the distance of the run, embodiments may share and/or offload certain tasks between the devices so that functionality will be available for the duration of the run. For example, if the battery power of the fitness tracker is low, GPS tracking can be offloaded to the phone device when a long run is anticipated. Alternatively, GPS tracking can be time division multiplexed between the two devices when a long run is anticipated to prolong battery power of one or more of the devices.

[0038] The following discussion now refers to a number of methods and method acts that may be performed. Although the method acts may be discussed in a certain order or illustrated in a flow chart as occurring in a particular order, no particular ordering is required unless specifically stated, or required because an act is dependent on another act being completed prior to the act being performed.

[0039] Referring now to Figure 2, a method 200 is illustrated. The method 200 includes acts that may be practiced in a computing environment. The method 200 includes acts for pooling computing resources based on inferences about a plurality hardware devices.

[0040] The method 200 includes identify inference information about the plurality of devices (act 202). For example, as demonstrated above, the inference information may be included in, or derived from device models 120. The inference information may include information that can be used to infer when devices will be proximate each other. The inference information may include information that can be used to infer when devices will be performing computing tasks. The inference information may include information that can be used to infer when devices will have spare resources available. Such resources may be computing resources, such as spare CPU cycles. Such resources may be network resources, such as network hardware bandwidth that is not being used. Such resources may

be storage resources, such as spare memory resources or mass storage resources that are available to be used. Such resources may be peripheral resources, such as various sensors that are not being used. Such resources may be applications that are available on devices.

5 [0041] The method 200 further includes, based on the inference information optimizing resource usage of the plurality of hardware devices (act 204). Various optimization examples will be illustrated below.

10 [0042] In one illustrative example, the inference information includes information identifying that the plurality of devices are likely to be proximate each within a given time. In one such embodiment, optimizing resource usage of the plurality of hardware devices may include delaying a computing operation in anticipation of the devices being in proximity to each other. For example, if embodiments know that a user is likely to have a phone device, and that the user and the phone device are likely to be proximate a laptop device or a desktop device at some time in the near future (where ‘near future’ is within some predetermined time), then computing on the phone device may be delayed until the
15 devices are proximate each other. This could allow for processing to be offloaded from the phone device to the laptop device or desktop device making the phone device seem more efficient to the user.

[0043] In another example, the plurality of hardware devices may include a laptop device (or other mobile device) and a power supply. If the inference information indicates
20 that a laptop device is likely to be coupled to a power supply in the near future, some computing tasks may be delayed until the laptop device is coupled to the power supply. This may allow the laptop device to be operated at higher computing power modes or may allow embodiments to delay processing for some power intensive, battery draining task, knowing that the device is anticipated to be coupled to a power supply in the near future.

25 [0044] In an alternative example, the plurality of devices may include a computing device and a particular network router. If inference information indicated that a computing device and a particular network router (e.g., coupled to a high capacity network, or the router itself having the ability to transfer large amounts of data) embodiments could delay some network intensive operation until the computing device and the network router are
30 coupled together.

[0045] In some embodiments, the system may implement retry logic based on the inference information. For example, if a computing device is performing some operation that fails, the inference information may indicate that retrying the operation should be

delayed as a better set of resources for performing the operation will likely be available within some predetermined time.

5 [0046] In some embodiments, deadline information may be used in conjunction with the inference information to make hardware device optimizations. For example, if a computing system knows a deadline for performing some computing task from the deadline information, but also knows from the inference information that certain resources are likely to be available soon on one or more highly capable devices, the operation can be postponed to be offloaded to the one or more highly capable devices.

10 [0047] The method 500 may be practiced where the inference information includes information identifying that the plurality of devices are likely to be in proximity to each other for a given period of time. Thus, computing decisions and device optimizations can be made in anticipation of devices being proximate each other for at least some period of time. Alternatively, the inference information indicating that the plurality of devices are likely to be in proximity to each other for a given period of time may indicate that after
15 that time, it is likely that one or more of the devices will no longer be proximate each other.

[0048] Thus, for example, embodiments may prioritize certain computing tasks knowing that resources are likely to be lost after the given period of time. For example, consider an example where the system knows that a user is at home where their phone
20 device and desktop device are in close proximity. The desktop device may automatically download podcasts and transfer them to the phone device. If the system of agents know that the user is likely to leave for work by some time, then the system may prioritize, at the desktop device, downloading the podcast and transferring it to the phone device to ensure that the podcast is available to the user before they leave for work.

25 [0049] In another example optimizing resource usage of the plurality of hardware devices may be used in transactional computing. For example, in some embodiments, optimizing resource usage of the plurality of hardware devices comprises offloading a transaction from a first device to a second device in proximity to the first device because the system determines, based on the inference information, that the transaction is likely to
30 be completed prior to the first device and second device no longer being in close proximity to each other. Alternatively or additionally, embodiments may determine to not retry a transaction or suppress transaction retry logic on the second device because, based on the inference information, the first device and second device are not likely to be in proximity to each other for a sufficient amount of time and/or that hardware resources may be lost

after the given period of time for the second device to complete the retry of the transaction.

[0050] Alternatively or additionally, embodiments may use inference information to prioritize work. Priorities may be based on what operations need to be performed and what resources are available, how long they will likely be available, what resources are likely to be available, when they are likely to be available, etc.

[0051] Further, the methods may be practiced by a computer system including one or more processors and computer-readable media such as computer memory. In particular, the computer memory may store computer-executable instructions that when executed by one or more processors cause various functions to be performed, such as the acts recited in the embodiments.

[0052] Embodiments of the present invention may comprise or utilize a special purpose or general-purpose computer including computer hardware, as discussed in greater detail below. Embodiments within the scope of the present invention also include physical and other computer-readable media for carrying or storing computer-executable instructions and/or data structures. Such computer-readable media can be any available media that can be accessed by a general purpose or special purpose computer system. Computer-readable media that store computer-executable instructions are physical storage media. Computer-readable media that carry computer-executable instructions are transmission media. Thus, by way of example, and not limitation, embodiments of the invention can comprise at least two distinctly different kinds of computer-readable media: physical computer-readable storage media and transmission computer-readable media.

[0053] Physical computer-readable storage media includes RAM, ROM, EEPROM, CD-ROM or other optical disk storage (such as CDs, DVDs, etc.), magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store desired program code means in the form of computer-executable instructions or data structures and which can be accessed by a general purpose or special purpose computer.

[0054] A “network” is defined as one or more data links that enable the transport of electronic data between computer systems and/or modules and/or other electronic devices. When information is transferred or provided over a network or another communications connection (either hardwired, wireless, or a combination of hardwired or wireless) to a computer, the computer properly views the connection as a transmission medium. Transmission media can include a network and/or data links which can be used to carry out desired program code means in the form of computer-executable instructions or data

structures and which can be accessed by a general purpose or special purpose computer. Combinations of the above are also included within the scope of computer-readable media.

[0055] Further, upon reaching various computer system components, program code means in the form of computer-executable instructions or data structures can be transferred automatically from transmission computer-readable media to physical computer-readable storage media (or vice versa). For example, computer-executable instructions or data structures received over a network or data link can be buffered in RAM within a network interface module (e.g., a “NIC”), and then eventually transferred to computer system RAM and/or to less volatile computer-readable physical storage media at a computer system. Thus, computer-readable physical storage media can be included in computer system components that also (or even primarily) utilize transmission media.

[0056] Computer-executable instructions comprise, for example, instructions and data which cause a general purpose computer, special purpose computer, or special purpose processing device to perform a certain function or group of functions. The computer-executable instructions may be, for example, binaries, intermediate format instructions such as assembly language, or even source code. Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the described features or acts described above. Rather, the described features and acts are disclosed as example forms of implementing the claims.

[0057] Those skilled in the art will appreciate that the invention may be practiced in network computing environments with many types of computer system configurations, including, personal computers, desktop computers, laptop computers, message processors, hand-held devices, multi-processor systems, microprocessor-based or programmable consumer electronics, network PCs, minicomputers, mainframe computers, mobile telephones, PDAs, pagers, routers, switches, and the like. The invention may also be practiced in distributed system environments where local and remote computer systems, which are linked (either by hardwired data links, wireless data links, or by a combination of hardwired and wireless data links) through a network, both perform tasks. In a distributed system environment, program modules may be located in both local and remote memory storage devices.

[0058] Alternatively, or in addition, the functionally described herein can be performed, at least in part, by one or more hardware logic components. For example, and without limitation, illustrative types of hardware logic components that can be used

include Field-programmable Gate Arrays (FPGAs), Program-specific Integrated Circuits (ASICs), Program-specific Standard Products (ASSPs), System-on-a-chip systems (SOCs), Complex Programmable Logic Devices (CPLDs), etc.

5 **[0059]** The present invention may be embodied in other specific forms without departing from its spirit or characteristics. The described embodiments are to be considered in all respects only as illustrative and not restrictive. The scope of the invention is, therefore, indicated by the appended claims rather than by the foregoing description. All changes which come within the meaning and range of equivalency of the claims are to be embraced within their scope.

CLAIMS

1. A computing device configured to be pooled with other devices, the device comprising:

an agent;

a model coupled to the agent;

wherein the agent is configured to access the model and models from other agents on other devices;

wherein the agent is configured to obtain from the models inference information about the device and the other devices; and

wherein the agent is configured to optimize resource usage of the device using the inference information about the device and the other devices.

2. The computing device of claim 1, wherein the inference information comprises information identifying that the plurality of devices are likely to be in a predetermined proximity to each other within a given time and wherein optimizing resource usage of the device using the inference information about the device and the other devices comprises delaying a computing operation in anticipation of a plurality of devices being in the predetermined proximity to each other within the given time.

3. The computing device of claim 1, wherein optimizing resource usage of the device using the inference information about the device and the other devices comprises assigning computing tasks to a device with available resources.

4. The computing device of claim 1, wherein the inference information comprises information identifying that the plurality of devices are likely to be in proximity to each other for a given period of time.

5. The computing device of claim 4, wherein optimizing resource usage of the plurality of hardware devices comprises prioritizing computing tasks based on an inference that hardware resources may be lost after the given period of time.

6. The computing device of claim 4, wherein optimizing resource usage of the plurality of hardware devices comprises suppressing transaction retry logic on a hardware device based on an inference that hardware resources may be lost after the given period of time.

7. The computing device of claim 4, wherein optimizing resource usage of the plurality of hardware devices comprises offloading a transaction to a hardware device based on an inference that hardware resources are likely to be available or used for the

given period of time and based on information indicating that the transaction is likely to be committed or aborted within the given period of time.

8. In a computing environment a method of pooling computing resources based on inferences about a plurality of hardware devices, the method comprising:

identifying inference information about the plurality of devices; and

based on the inference information optimizing resource usage of the plurality of hardware devices.

9. The method of claim 8, wherein the inference information comprises information identifying that the plurality of devices are likely to be in a predetermined proximity to each other within a given time.

10. The method of claim 9, wherein optimizing resource usage of the plurality of hardware devices comprises delaying a computing operation in anticipation of the plurality of devices being in the predetermined proximity to each other within the given time.

11. The method of claim 9, wherein optimizing resource usage of the plurality of hardware devices comprises assigning computing tasks to a device with available resources

12. The method of claim 8, wherein the inference information comprises information identifying that the plurality of devices are likely to be in proximity to each other for a given period of time.

13. The method of claim 12, wherein optimizing resource usage of the plurality of hardware devices comprises prioritizing computing tasks based on an inference that hardware resources may be lost after the given period of time.

14. The method of claim 12, wherein optimizing resource usage of the plurality of hardware devices comprises suppressing transaction retry logic on a hardware device based on an inference that hardware resources may be lost after the given period of time.

15. The method of claim 12, wherein optimizing resource usage of the plurality of hardware devices comprises offloading a transaction to a hardware device based on an inference that hardware resources are likely to be available or used for the given period of time and based on information indicating that the transaction is likely to be committed or aborted within the given period of time.

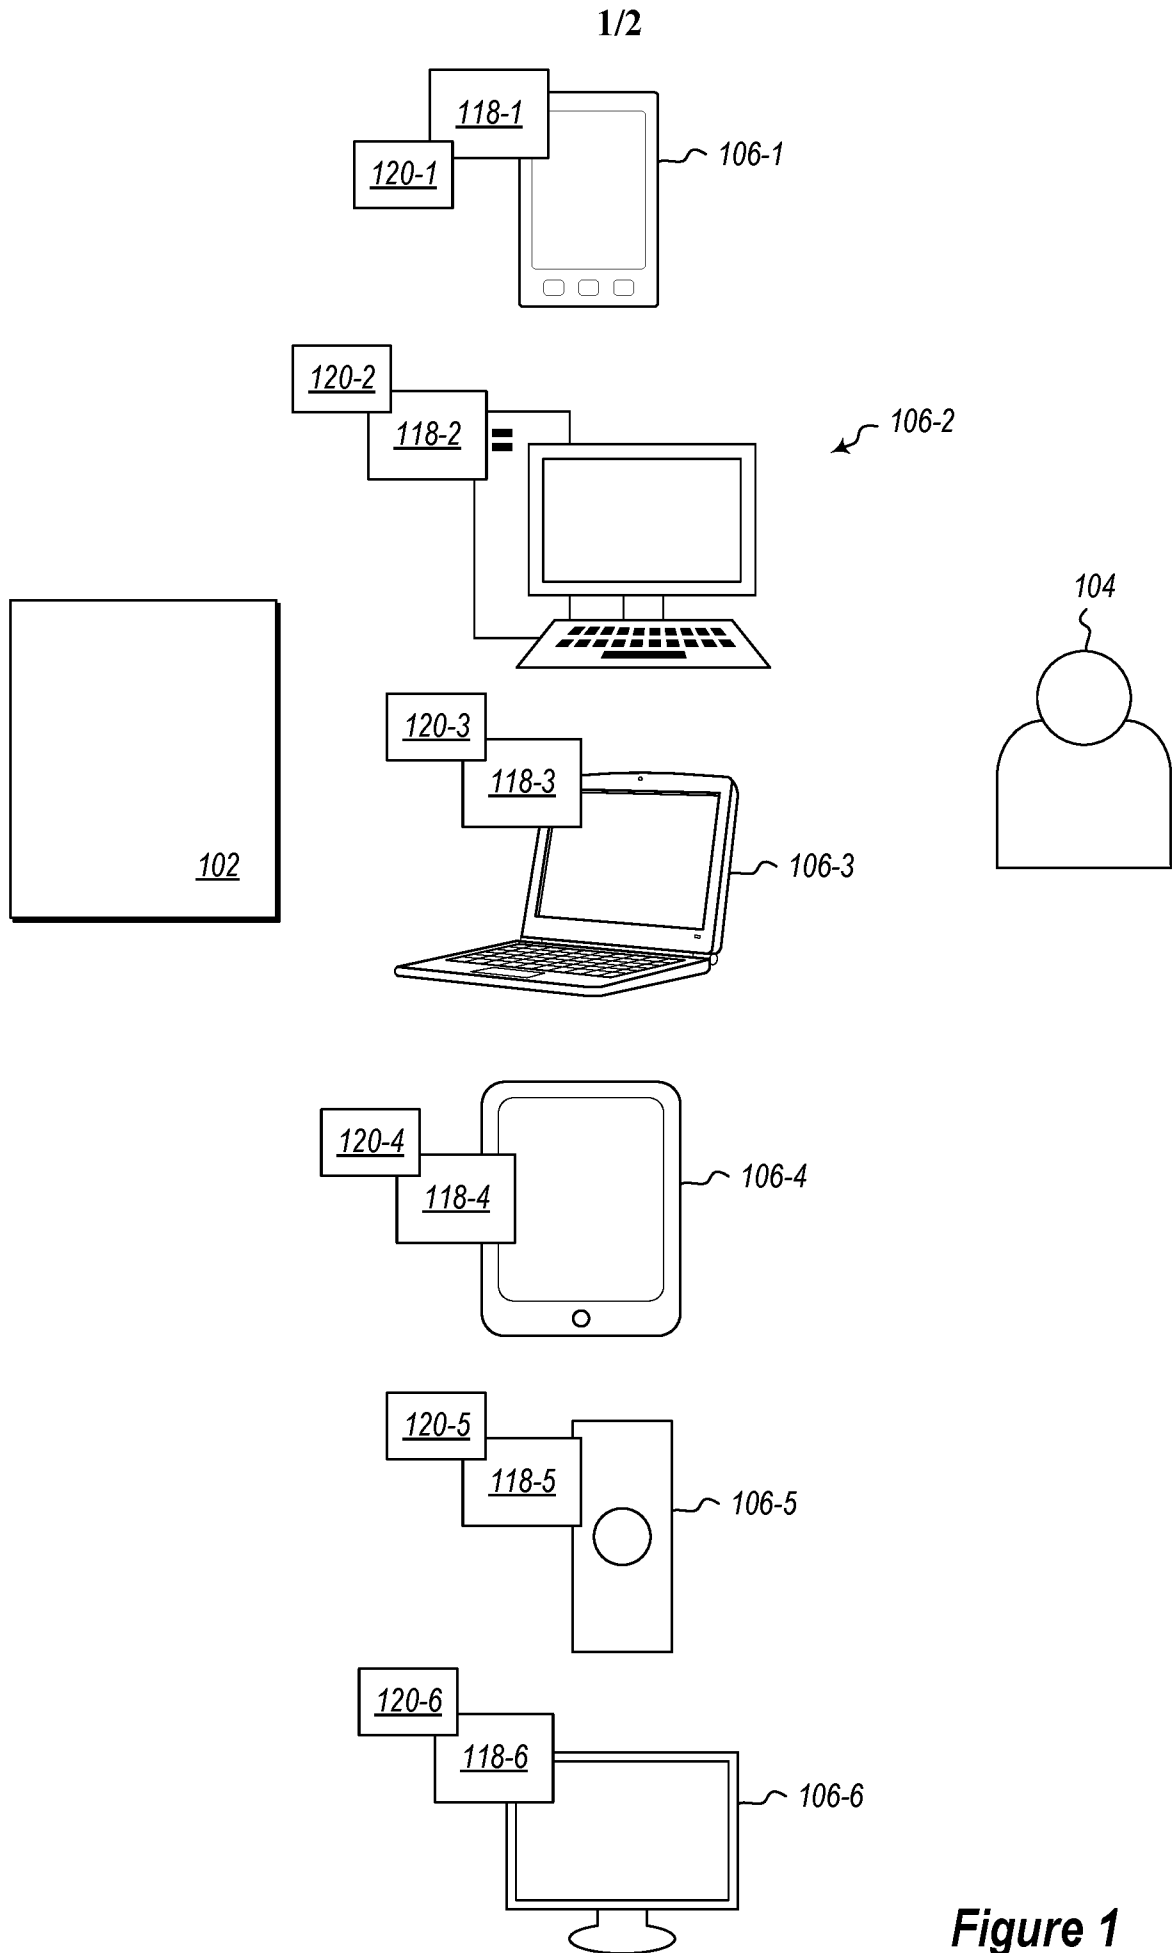


Figure 1

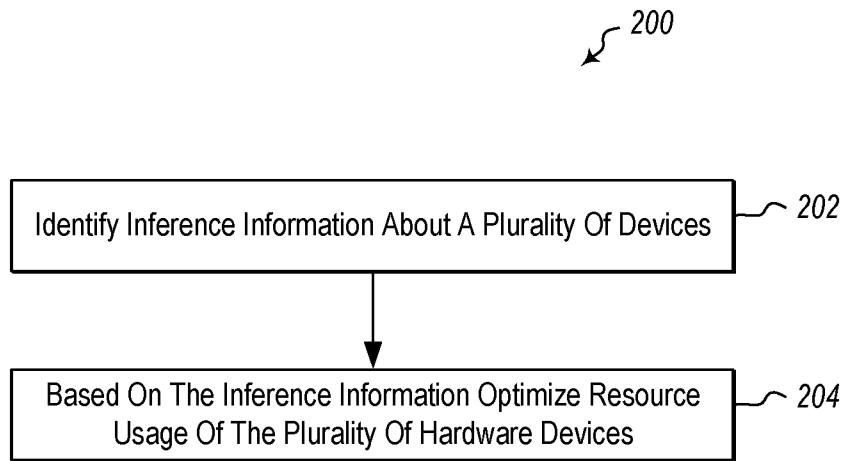


Figure 2

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2017/017090

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F9/50
ADD.
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED
Minimum documentation searched (classification system followed by classification symbols)
G06F

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

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
|-----------|--|-----------------------|
| X | US 2012/271953 A1 (GULLEY EDWARD WHITTINGTON [US] ET AL) 25 October 2012 (2012-10-25) paragraphs [0046], [0083], [0085], [0131] - [0132], [0141] - [0142], [0203], [0209], [0258] - [0282], [0310] ----- | 1-15 |
| X | WO 2008/107697 A1 (BRITISH TELECOMM [GB]; RINGLAND SIMON PATRICK A [GB]; BAPPU BENJAMIN []) 12 September 2008 (2008-09-12) page 2, line 20 - line 32 page 3 page 4, line 15 - line 33 page 5, line 1 - line 8 page 19, line 19 - line 32 page 23, line 27 - line 33 page 24, line 1 - line 7 ----- -/-- | 1-15 |

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

| | |
|---|--|
| "A" document defining the general state of the art which is not considered to be of particular relevance | "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention |
| "E" earlier application or patent but published on or after the international filing date | "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone |
| "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) | "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art |
| "O" document referring to an oral disclosure, use, exhibition or other means | "&" document member of the same patent family |
| "P" document published prior to the international filing date but later than the priority date claimed | |

| | |
|---|--|
| Date of the actual completion of the international search 4 May 2017 | Date of mailing of the international search report 18/05/2017 |
|---|--|

| | |
|--|---------------------------------|
| Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016 | Authorized officer Simion, C |
|--|---------------------------------|

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2017/017090

| C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT | | |
|--|---|-----------------------|
| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
| X | US 2008/320482 A1 (DAWSON CHRISTOPHER J [US] ET AL) 25 December 2008 (2008-12-25) paragraphs [0009] - [0011], [0021], [0022], [0029], [0031], [0033] - [0035], [0044], [0045] ----- | 1-15 |
| X | US 9 195 508 B1 (BLANDING WILLIAM H [US] ET AL) 24 November 2015 (2015-11-24) column 1, line 50 - line 67 column 2 column 3, line 1 - line 31 column 5, line 61 - line 67 column 6, line 1 - line 22 ----- | 1-15 |
| A | US 2013/262644 A1 (HINTERMEISTER GREGORY R [US] ET AL) 3 October 2013 (2013-10-03) paragraphs [0006], [0018], [0031], [0033] ----- | 2,10 |

INTERNATIONAL SEARCH REPORT

Information on patent family members

| |
|---|
| International application No PCT/US2017/017090 |
|---|

| Patent document cited in search report | Publication date | Patent family member(s) | Publication date |
|--|------------------|-------------------------|------------------|
| US 2012271953 | A1 | 25-10-2012 | EP 2504764 A2 |
| | | | US 2010223385 A1 |
| | | | US 2012271939 A1 |
| | | | US 2012271953 A1 |
| | | | US 2015106498 A1 |
| | | | WO 2011081757 A2 |
| ----- | | | |
| WO 2008107697 | A1 | 12-09-2008 | EP 1968243 A1 |
| | | | EP 2119122 A1 |
| | | | US 2010085947 A1 |
| | | | WO 2008107697 A1 |
| ----- | | | |
| US 2008320482 | A1 | 25-12-2008 | TW 200915186 A |
| | | | US 2008320482 A1 |
| ----- | | | |
| US 9195508 | B1 | 24-11-2015 | NONE |
| ----- | | | |
| US 2013262644 | A1 | 03-10-2013 | NONE |
| ----- | | | |