



(51) International Patent Classification:

G06F 9/50 (2006.01) G06F 11/34 (2006.01)

(21) International Application Number:

PCT/US2019/058526

(22) International Filing Date:

29 October 2019 (29.10.2019)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

16/220,419 14 December 2018 (14.12.2018) US

(71) Applicant: **UBER TECHNOLOGIES, INC.** [US/US];
1455 Market St., 4th Floor, San Francisco, California 94103
(US).(72) Inventor: **KAO, Shao-Hua**; 1455 Market St., 4th Floor,
San Francisco, California 94103 (US).(74) Agent: **JONSSON, Jon G.**; Mahamedi IP Law LLP, 910
Campisi Way, Suite 1E, Campbell, California 95008 (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, JO, 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, 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.(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, SI, SK, SM,
TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
KM, ML, MR, NE, SN, TD, TG).

Published:

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

(54) Title: MEMORY CRASH PREVENTION FOR A COMPUTING DEVICE




FIG. 3

(57) Abstract: A computing device can monitor a set of memory usage metrics of the computing device. Based on historical memory usage data and the set of memory usage metrics, the computing device can determine whether memory usage will exceed a critical memory threshold at a future instance in time. In response to determining that the memory usage will exceed the critical memory threshold at the future instance in time, the computing device can degrade one or more application features of an application executing on the computing device.

MEMORY CRASH PREVENTION FOR A COMPUTING DEVICE

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of priority to U.S. Patent Application No. 16/220,419, filed December 14, 2018; the aforementioned priority application being hereby incorporated by reference in its entirety.

BACKGROUND

[0002] Out of memory (OOM) events occur on computing devices when little or no additional memory can be allocated for use by the operating system, applications, or other programs. The result of such a condition can cause applications to cease functioning properly or cause the application to crash.

BRIEF DESCRIPTION OF THE DRAWINGS

[0003] The disclosure herein is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings in which like reference numerals refer to similar elements, and in which:

[0004] FIG. 1 is a block diagram illustrating an example computing device implementing memory management and crash prevention techniques, in accordance with examples described herein;

[0005] FIG. 2 depicts an example illustration of a hierarchical table of application features, according to examples described herein;

[0006] FIGS. 3 and 4 are block diagrams illustrating example methods of managing memory usage for an executing application on a computing device, according to examples described herein; and

[0007] FIG. 5 is a block diagram that illustrates a computing device upon which examples described herein may be implemented.

DETAILED DESCRIPTION

[0008] A computing device can monitor a set of memory usage metrics of the computing device. For example, the computing device may be executing one or more applications that include or display features or perform operations requiring extensive memory usage. In some examples, based on historical memory usage data and the set of memory usage metrics, an application running on the computing device (or a system or

operating system of the computing device, in another implementation) (referred to herein as the computing device, for purposes of simplicity) can determine or predict whether future memory usage will exceed a critical memory threshold. In response to determining that the future memory usage will exceed the critical memory threshold, the computing device can degrade and/or deactivate one or more application features of the application.

[0009] In various implementations, the executing application can comprise a set of features, each of which can be associated with one or more hierarchical tier levels. For example, a particular application feature can be essential to the functioning of the application, and can be assigned a highest tier. The highest tiered application features can be protected from deactivation or suspension, but as described below, may be classified as available for degradation. In certain examples, these highest tier application features can be excluded from degradation and/or deactivation by the computing device. Other features can be classified in one or more non-essential tiers, making such features available for degradation and/or deactivation.

[0010] According to examples described herein, the application features can further be classified into subcategories (e.g., as a metadata tag for each feature) indicating whether the feature is available for deactivation, degradation and deactivation, degradation only (e.g., and levels of degradation), neither deactivation nor degradation. As used herein, “degradation” of an application feature can comprise reducing a content render rate of the feature (e.g., for streaming or updating content), expanding tiles (e.g., reducing display resolution by a target percentage), reducing a data transfer rate (e.g., increasing location data updating periods), and the like. As further provided herein, “deactivation” of an application feature can comprise eliminating the feature from current usage (or suspend the feature from current usage for a duration of time) on the application to free up the current memory usage of that feature.

[0011] In various examples, the computing device can include memory monitoring functionality that monitors a set of memory metrics of the computing device. In one aspect, the computing device dynamically monitors a current memory usage by the executing application. The computing device can further determine the system free memory of the computing device, and/or the amount of total memory the computing device contains. In further implementations, the computing device can dynamically determine a ratio of the total free memory versus current memory usage of the executing application. In still further implementations, the computing device can store historical data indicating historical memory usage for the executing application, and monitor a

predetermined period of current memory usage (e.g., the previous five seconds of memory usage by the application).

[0012] According to examples described herein, the computing device can project a future memory usage for the executing application based on the predetermined period of current memory usage and the historical memory usage of the application. For example, the computing device can determine whether the executing application will exceed a critical memory usage threshold within the next thirty seconds. Additionally or alternatively, the computing device can leverage third-party memory warning systems, such as operating system (OS) integrated memory warning systems included with the computing device. Such third-party warning systems can provide low memory warnings that cause the computing device to perform a first set of memory usage reduction actions, such as closing background applications not currently in use. In addition, the computing device can log each memory warning from the OS memory warning system to determine a memory pressure level for the executing application.

[0013] In various examples, threshold numbers of low memory warnings from the OS memory warning system over a given time can be established by the computing device as tiered memory pressure levels for the executing application. For example, three memory warnings within the previous one minute can comprise a high memory pressure level, and five memory warnings within the previous minute can comprise a critical memory pressure level. Each memory pressure level can trigger the computing device to perform a respective set of memory usage reduction actions for the executing application. For example, when the high memory pressure threshold is met or exceeded, the computing device can begin to degrade and/or deactivate a lower tiered set of application features (e.g., non-essential features). If the critical memory pressure threshold is met, the computing device can aggressively degrade and/or deactivate additional application features (e.g., higher tiered application features) to prevent an out-of-memory (OOM) state for the computing device.

[0014] Additionally or alternatively, the computing device can predict a future memory usage based on historical application usage data and a brief period of application usage (e.g., the previous five seconds), and determine whether the predicted future memory usage will exceed a particular memory pressure level. Thus, the tiered memory reduction triggers can be based on the prediction by the computing device. For example, based on a memory usage projection from historical data and current usage data, the computing device can predict that a third-tier memory pressure level will be exceeded in

the next twenty seconds. In response to this prediction, the computing device can preemptively implement memory usage reduction actions to prevent the third-tier memory threshold level from being exceeded in the first place. As an example, the computing device can perform a lookup of third-tier application features and the corresponding memory reduction actions to perform on such features, and then preemptively implement the memory reduction actions accordingly.

[0015] In various implementations, the computing device can monitor the memory metrics to dynamically determine whether a first-tier memory pressure threshold is exceeded. If so, the computing device can perform first-tiered memory usage reduction actions for an executing application. The computing device can further determine whether a second-tier, more critical, memory pressure threshold is exceeded. If so, the computing can perform second-tiered memory usage reduction actions for the executing application. It is contemplated that the computing device can establish any number of tiers for memory pressure levels, and can further perform any number of tiers of corresponding memory usage reduction actions. It is further contemplated that each tier of memory pressure level can correspond to a particular tier of memory usage reduction actions. Furthermore, each tier of memory usage reduction actions can correspond to a set of application features, which can be selectively degraded or deactivated in response to the corresponding memory pressure level being exceeded.

[0016] As an example, the executing application can comprise a navigation application comprising highest tiered features such as live map content, a guided trajectory overlaying the map content, updates to the trajectory, and the like. Certain features of the highest tier can be excluded from degradation, such as the guided trajectory feature, where the map content may be available for degradation (e.g., reduction of content rendering). The navigation application features can further include non-essential features, such as voiced instructions to prepare for upcoming turns and to follow the guided trajectory, traffic alerts, overlaid indicators for fueling or charging stations, restaurants, points of interest, and non-essential buttons or selectable features. In one example, the application can further include non-essential, memory-intensive features, such as streaming or near-continuously updating data (e.g., surge pricing regional indicators for transport service applications utilized by transport providers).

[0017] As the computing device receives memory warnings, a first memory pressure level may be exceeded. In response to the first memory pressure level being exceeded, the computing device can identify a first set of application features of the navigation

application, and a set of memory reduction actions to perform on the first set of application features. For example, the memory reduction action for each application feature of the first set of application features can be identified in metadata as either degradation or deactivation, and can further define the manner in which the feature is to be degraded (e.g., reduce a map content render rate by 20%).

[0018] Once this first set of memory reduction actions has been performed on the executing application, the computing device can continue to monitor the memory metrics, and/or monitor warnings issued by the OS memory warning system. If a second memory pressure threshold is exceeded, the computing device can determine a second set of application features, and a set of memory reduction actions to perform on the second set of application features. It is to be noted that the second set of application features can include application features from the first set (e.g., as degraded features). Furthermore, the memory reduction action for each of the second set of application features can be determined and implemented by the computing device. These actions can include degradation, further degradation (e.g., for features in the first set as well as the second set), and deactivation.

[0019] The computing device can continue to monitor the memory metrics for additional memory pressure thresholds being exceeded. Accordingly, for each exceeded memory pressure threshold, the computing device can look up a corresponding set of application features for the executing application, and also a corresponding set of memory reduction actions to perform on those application features. Thus, for the executing application, each feature can be listed in a database or logical table, each feature can be associated with one or more tiers for memory reduction, and for each listed tier, each feature can be associated with a memory reduction action (e.g., degradation, degradation level, or deactivation).

[0020] In one example, the executing application can comprise a transport service application utilized by transport providers of a network-based transport service for a given region. The transport service application can include various memory-intensive features that require continuous or near continuous updating. Such features include map content, transport invitation notifications, timers for accepting or declining transport invitations, navigation content (e.g., a turn-by-turn trajectory overlaying the map content), voiced turn-by-turn instructions, selectable buttons or menu items, selectable map features (e.g., information content for points of interest), and the like. The features can further include a dynamic pricing feature that provides the transport provider with

information on map content that indicates high demand areas for the transport service. This dynamic pricing feature requires near continuous updating based on supply-demand data for the transport service, and further includes rendering overlaid content on the map data (e.g., color-coded shading over sub-regions having high transport demand, a surge pricing multiplier feature, etc.).

[0021] In such an example, the computing device can store a tiered table listing each tier of application features for the transport service application. Furthermore, the table can further list corresponding memory pressure levels for each application feature, and a corresponding memory reduction action to be performed when that memory pressure threshold is exceeded. For example, the dynamic pricing feature of the transport service application can be associated with multiple tiers, each succeeding tier indicating degradation actions to be performed, with a final tier indicating deactivation of the dynamic pricing feature. When a first memory pressure threshold level is exceeded, the computing device can perform a lookup in the table, identify the dynamic pricing feature in the first set of application features, and determine a memory reduction action to perform on the dynamic pricing feature (e.g., reduce a content render rate). When a n th memory pressure threshold level is exceeded, the computing device can perform a lookup in the table for the n th tiered set of application features and the corresponding memory reduction actions to perform on those application features.

[0022] Among other benefits, examples described herein achieve a technical effect of reactively and/or preemptively resolving low or out-of-memory conditions for a computing device by implementing highly granular application feature degradation for a target application executing on the computing device. Such examples can implement predictive functionality as well as leveraging current low memory warning techniques to optimize the performance of the target application given current or predicted memory conditions.

[0023] As used herein, a computing device refers to devices corresponding to desktop computers, cellular computing devices and/or smartphones, personal digital assistants (PDAs), laptop computers, virtual reality (VR) or augmented reality (AR) headsets, tablet computing devices, etc., that can provide network connectivity and processing resources for communicating with the system over a network. A computing device can also correspond to custom hardware, in-vehicle devices of automobiles, or on-board computers, etc. The computing device can also operate a designated application configured to communicate with the network service.

[0024] One or more examples described herein provide that methods, techniques, and actions performed by a computing device are performed programmatically, or as a computer-implemented method. Programmatically, as used herein, means through the use of code or computer-executable instructions. These instructions can be stored in one or more memory resources of the computing device. A programmatically performed step may or may not be automatic.

[0025] One or more examples described herein can be implemented using programmatic modules, engines, or components. A programmatic module, engine, or component can include a program, a sub-routine, a portion of a program, or a software component or a hardware component capable of performing one or more stated tasks or functions. As used herein, a module or component can exist on a hardware component independently of other modules or components. Alternatively, a module or component can be a shared element or process of other modules, programs or machines.

[0026] Some examples described herein can generally require the use of computing devices, including processing and memory resources. For example, one or more examples described herein may be implemented, in whole or in part, on computing devices such as servers, desktop computers, cellular or smartphones, personal digital assistants (e.g., PDAs), laptop computers, VR or AR devices, printers, digital picture frames, network equipment (e.g., routers) and tablet devices. Memory, processing, and network resources may all be used in connection with the establishment, use, or performance of any example described herein (including with the performance of any method or with the implementation of any system).

[0027] Furthermore, one or more examples described herein may be implemented through the use of instructions that are executable by one or more processors. These instructions may be carried on a computer-readable medium. Machines shown or described with figures below provide examples of processing resources and computer-readable mediums on which instructions for implementing examples disclosed herein can be carried and/or executed. In particular, the numerous machines shown with examples of the invention include processors and various forms of memory for holding data and instructions. Examples of computer-readable mediums include permanent memory storage devices, such as hard drives on personal computers or servers. Other examples of computer storage mediums include portable storage units, such as CD or DVD units, flash memory (such as carried on smartphones, multifunctional devices or tablets), and magnetic memory. Computers, terminals, network enabled devices (e.g., mobile devices,

such as cell phones) are all examples of machines and devices that utilize processors, memory, and instructions stored on computer-readable mediums. Additionally, examples may be implemented in the form of computer-programs, or a computer usable carrier medium capable of carrying such a program.

[0028] SYSTEM DESCRIPTION

[0029] FIG. 1 is a block diagram illustrating an example computing device 100 implementing memory management and crash prevention techniques, in accordance with examples described herein. The computing device 100 can include a set of sensors 150 (e.g., positioning system, IMU, camera, microphone, touch sensors, ambient light sensor, etc.) and a display 160 generating a display interface 162 for an executing application 135. The computing device 100 can further include a communication interface 110 that connects the computing device 100 with one or more networks 180. The computing device 100 can further include a memory monitor 120 and application manager 140. The memory monitor 120 can continuously or periodically monitor a set of memory metrics (e.g., current usage data 132 and/or historical usage data 134) for memory 130 of the computing device 100.

[0030] The application manager 140 can execute applications 135 in response to input data from a user. The memory 130 can store a set of applications 135. For certain applications 135, the memory 130 can further store an application feature hierarchy table 136 listing each application feature in a tiered hierarchy corresponding to memory pressure threshold levels. The application feature hierarchy table 136 can include memory usage reduction action metadata 138 indicating the memory reduction actions to be performed for each application feature at each hierarchical tier (e.g., no action, degradation, degradation level, deactivation).

[0031] In various implementations, the memory 130 can collect and store historical memory usage data 134 specific to the executing application 135. For example, the historical usage data 134 can indicate the memory usage of the executing application for a predetermined period of time prior to the present (e.g., thirty seconds). In contrast, the current usage data 132 can correspond to a shorter amount of time prior to the present (e.g., five seconds). Based on the historical memory usage data 134 and the current usage data 132 for the executing application, the memory monitor 120 can determine whether memory usage for the computing device 100 will exceed a critical memory threshold at a future instance in time (e.g., thirty seconds in the future). As provided herein, the critical

memory threshold can correspond to an impending out-of-memory crash on the computing device 100.

[0032] The memory monitor 120 can utilize the memory metrics to predict whether a particular memory pressure threshold will be exceeded, and issue warning triggers—indicating a current or predicted memory threshold being exceeded—to the application manager 140 accordingly. In various examples, the memory metrics can comprise the current system free memory of the computing device 100, and the amount of total memory that the computing device 100 contains. In further implementations, the computing device 100 can dynamically determine or otherwise predict a ratio of the total free memory versus current memory usage of the computing device 100. If the current or predicted ratio exceeds a first threshold (e.g., a first warning trigger by the memory monitor 120), then the application manager 140 can perform a lookup in the hierarchy table 136 for the first set of features on which to implement memory reduction actions. If the ratio exceeds a second threshold, application manager 140 can perform a second lookup in the hierarchy table 136 for a second set of application features on which to implement memory reduction actions, and so on and so forth down the hierarchy. Accordingly, based on the level of warning trigger from the memory monitor 120, the application manager 140 can perform a lookup in the application feature hierarchy table 136 of the executing application for a corresponding set of application features, and then implement the set of memory usage reduction actions indicated in the memory usage reduction action metadata 138 for those features.

[0033] As described herein, the memory monitor 120 can also leverage the functions of an OS memory warning system 131 of the computing device 100, which can issue reactive warnings to the memory monitor 120 based on the current memory conditions. In certain aspects, the memory monitor 120 can utilize these warnings to determine a memory pressure warning level for the executing application 135. In variations, the memory monitor 120 can utilize the reactive warnings from the OS warning system 131 to predict whether the computing device 100 will exceed a particular memory pressure threshold of multiple memory pressure thresholds at a future instance in time (e.g., thirty seconds in the future). In doing so, the memory monitor 120 can determine a number of reactive memory warnings from the OS memory warning system 131 for the computing device 100 as a whole, the current usage data 132 for the target executing application 135, and historical usage data 134 for that application. Based on these factors, the memory monitor 120 can predict or otherwise determine a memory pressure threshold level—of

the multiple memory pressure levels—that will be exceeded, and submit a corresponding warning trigger to the application manager 140.

[0034] In response to a warning trigger from the memory monitor 120 indicating that the memory usage will exceed a particular memory threshold at the future instance in time (e.g., a critical memory threshold indicating an imminent crash), the application manager 140 can degrade one or more application features of the executing application 135. For example, the application manager 140 can perform a lookup in the application feature hierarchy table 136 of the executing application 135 to identify a set of application features for degradation and/or deactivation (e.g., a tiered set of features that correspond to the memory pressure threshold determined by the memory monitor 120). In certain implementations, the hierarchy table 136 can include a listing of application features and usage reduction metadata 138 tagged to each feature to indicate which particular memory reduction action to implement for each application feature.

[0035] It is contemplated that the application manager 140 can comprise a component of the executing application 135 itself, or can comprise a component of the operating system executing on the computing device 100. Thus, the functions of the application manager 140 described herein can be configured by the operating system to monitor and manage memory of any application executing on the computing device 100, or can comprise a monitoring and management component of any individual application executing on the computing device.

[0036] As described herein, the executing application 135 in which the application manager 140 focuses degradation, can comprise a content streaming application (e.g., video or music streaming), a mapping and/or navigation application, a social media application, a news content application, a marketplace or sales application, and the like. In certain examples, the application manager 140 can comprise a built-in management program of the executing application 135, enabling the application developer to create the application feature hierarchy table 136. Such an arrangement enables the application developer to prioritize which features are essential, which features are essential but may be degraded, and which features are not essential and may be degraded and/or deactivated.

[0037] In certain aspects, the memory monitor 120 can issue hierarchical warnings to the application manager 140. For example, a first hierarchical warning can comprise a lower tiered memory usage warning, and can cause the application manager 140 to only degrade a first set of non-essential application features, and/or deactivate certain non-

essential features. A next hierarchical warning from the memory monitor 120 can cause the application manager 140 to lookup a next hierarchical level of application features in the hierarchy table, and perform a set of memory reduction action on those features (e.g., as indicated in the usage reduction metadata 138 for those features within that hierarchical tier). A critical memory warning (e.g., indicating an impending out-of-memory crash) from the memory monitor 120 can cause the application manager 140 to look up a last tiered set of application features in the hierarchy table 136 and perform a final, most aggressive set of memory usage reduction actions on those features.

[0038] Examples of memory reduction actions can comprise reduction of resolution of content data to be presented on a display interface 162 of the display screen 160 (e.g., expanding or joining pixels or display tiles). Such actions can further include reducing a content render rate of a content stream or mapping content (e.g., from ten frames per second to five frames per second). Additionally or alternatively, the memory reduction actions can include reducing a sensor data transmission rate from the sensors 150 (e.g., transmitting position updates every eight seconds as opposed to every four seconds), or deactivate of certain sensors 150. In certain examples, the application manager 140 can also reduce a service data update rate received from a datacenter of the networks 180. For example, under normal memory conditions, the service data may be updated every ten seconds. Such updates may indicate traffic condition updates, supply-demand updates for transport services, news feed updates, and the like. When a particular memory pressure threshold is exceeded, the application manager 140 can decrease the service data update rate to, for example, once every thirty seconds.

[0039] It is contemplated that such granular feature management for an application 135 can result in both the optimization of that application given current and/or predicted memory conditions of the computing device 100, and the prevention of an out-of-memory state for the computing device 100, which can cause multiple applications to function improperly. For example, memory-intensive applications—such as social media, navigation, and content streaming—can leverage the use of the memory monitor 120 and application manager 140 to significantly reduce the memory consumption of a target application 135 to alleviate high or critical memory conditions. The application 135 can include built-in application management functionality, which can utilize an application feature hierarchy table 136 and reduction metadata 138 described throughout the present disclosure to implement memory usage reduction techniques described herein.

[0040] HIERARCHICAL APPLICATION FEATURE TABLE

[0041] FIG. 2 depicts an example illustration of a hierarchical table 200 of application features, according to examples described herein. In the below description of FIG. 2, reference may be made to reference characters representing like features as shown and described with respect to FIG. 1. Referring to FIG. 2, the hierarchical table 200 can be stored in a database 295 of the computing device 100 and can correspond to the hierarchical table 136 described with respect to FIG. 1. In certain aspects, the hierarchical table 200 can be associated with a single application of the computing device 100, such as a navigation application, a transport service application, a content streaming application, and the like. In certain implementations, the hierarchical table 200 can include hierarchical feature tiers for multiple device applications 205, as shown in FIG. 2. For example, any of the listed applications 205 may be currently executing, or can be running as a background application when a certain memory pressure threshold is exceeded. In response to the threshold being exceeded, the computing device 100 can identify any of the currently executing applications 205, and perform lookups for respective tiered application features for that particular application.

[0042] As shown in FIG. 2, an executing application (App 1) can comprise numerous application features listed in a feature hierarchy 245 for memory reduction treatment when specified memory pressure thresholds are exceeded. In accordance with examples described herein, the memory monitor 120 can transmit a first memory pressure warning trigger 212 to the application manager 140. In response, the application manager 140 can perform a first lookup in the feature hierarchy table 200 for the executing application (App 1) to identify first tier application features 210 for memory usage reduction. In certain aspects, each of the first tier application features 210 can be associated with a memory reduction action 215 (e.g., in metadata) for the application manager 140 to perform.

[0043] With reference to FIG. 2, App 1 can comprise first tier features 210 comprising voice commands, a menu list, map content, GPS data transmission to a remote source, a dynamic pricing feature (e.g., for transport service applications), and an advertisement updating feature. The action tags 215 for each feature can instruct the application manager 140 to perform a specified memory reduction action on the feature. In the example shown, the memory reduction actions 215 comprise deactivation of the voice commands, a reduction of the menu list (e.g., eliminating low priority items), degradation of the map content (e.g., by 20%), degradation of GPS data transmission rate (e.g.,

reduced from once every four second to once every eight seconds), degradation of the dynamic pricing updates (e.g., by 50%), and degradation of ad updates (e.g., by 50%).

[0044] In response to a second, more critical memory pressure trigger 222, the application manager 140 can perform a lookup of second tier application features 220 and implement the respective memory reduction actions associated with each of the second tier application feature 220. It is to be noted that some of the first tier application features 210 are also listed in the second tier 220, but are associated a different, more aggressive memory reduction action. Furthermore, it is contemplated that the application manager 140 need not proceed down the feature hierarchy 245 sequentially. Rather, the memory monitor 120 may issue any memory pressure level warning 212, 222, corresponding to any tier of the hierarchical table 200 at any given time, and the application manager 140 can perform each of the memory reduction actions (e.g., deactivation actions) on any lower-tiered application features, and then perform the memory reduction actions on application features corresponding to the current memory pressure level warning.

[0045] For example, the memory monitor 120 may issue a highly critical, third level memory pressure warning, corresponding to the third tier application features 230 of App 1. In response, the application manager 140 can deactivate voice commands from the first tier 210, then deactivate the ad updates, menu list, dynamic pricing feature, and the point of interest selection feature from the second tier 220, and then finally perform the memory reduction actions on the third tier application features 230 (e.g., degrade the trajectory line by 10%, deactivate the content stream, and degrade the map content by 60%). As described herein, the feature hierarchy 240 can comprise any number of tiers, each corresponding to a particular memory pressure threshold, or memory pressure level trigger 212, 22, as shown in FIG. 2. It is contemplated that this hierarchical memory reduction treatment of a specified application can prevent out-of-memory crashes while optimizing the application given the current memory conditions of the computing device 100.

[0046] METHODOLOGY

[0047] FIGS. 3 and 4 are block diagrams illustrating example methods of managing memory usage for an executing application on a computing device, according to examples described herein. In the below descriptions of FIGS. 3 and 4, reference may be made to reference characters representing like features as shown and described with respect to FIGS. 1 and 2. Furthermore, the processes described in connection with FIGS. 3 and 4 may be performed by an example computing device 100 as shown and described

with respect to FIG. 1. Referring to FIG. 3, the computing device 100 can monitor a set of memory usage metrics of the computing device 100 (300). As described herein, the memory usage metrics can comprise a the size of memory currently used by an executing application, the total system free memory of the computing device 100, and/or the system free memory to total system size ratio.

[0048] In various implementations, the computing device 100 can analyze historical memory usage data 134 of the computing device 100 (305). Specifically, the computing device 100 can analyze the historical memory usage data of the executing application (307). In certain aspects, the historical data 134 can comprise previous application usage sessions by the user, or can comprise a set amount of time prior to the present (e.g., the past N seconds) (309). Based on the set of memory usage metrics and the historical data 134, the computing device 100 can determine or predict whether a critical memory threshold will be exceeded (e.g., in the next twenty seconds) (310). For example, the computing device 100 can utilize the historical memory usage data 134, current usage data 132 and the current memory metrics to predict whether the critical memory threshold will be exceeded at a future instance in time. It is contemplated that such a determination can comprise a probabilistic calculation, and a “yes” can comprise a probability above a certain threshold (e.g., 90%), whereas a “no” can be anything below the threshold.

[0049] If the computing device determines that the critical memory threshold will not be exceeded (312), then the computing device 100 can continue to monitor the set of memory usage metrics accordingly (300). However, if the computing device determines that the critical memory threshold will be exceeded (314), then the computing device 100 can degrade and/or deactivate certain features of the executing application (315). For example, the computing device 100 can reduce a content render rate for the application (316), shut down one or more memory intensive features (317), and/or reduce a resolution of content on the display 160 (318).

[0050] FIG. 4 is a flow chart describing another example method of managing memory usage of an executing application given current and/or historical memory usage data, according to various examples. Referring to FIG. 4, the computing device 100 can determine a system free memory of the computing device 100 (400). The computing device 100 can further periodically or continuously determine a ratio of system free memory to total memory for the computing device 100 (405). The computing device 100 can monitor the ratio based on a set of hierarchical memory pressure level thresholds (410). In one example, the computing device 100 can monitor a ratio of the amount of

memory being used by the executing application versus the system free memory available on the computing device 100. Additionally or alternatively, the computing device 100 can include an OS memory usage warning system, which can be utilized to predict whether one or more memory pressure thresholds will be exceeded. Based at least in part on the memory usage ratio, the computing device 100 can determine or predict whether a first memory pressure threshold will be exceeded at a future instance in time (415). If not (417), then the computing device 100 can continue to monitor the ratio and/or other memory metrics accordingly (410).

[0051] However, if a first memory pressure threshold is predicted to be exceeded (419), then the computing device 100 can degrade and/or deactivate a set of first tier application features, as described herein (430). The computing device 100 may then continue to monitor the memory metrics and/or memory usage ratio (410). The computing device 100 may determine whether a next memory pressure threshold will be exceeded at a future instance (415). If not (417), then the computing device 100 will continue monitoring (410). However, if so (429), then the computing device 100 can degrade and/or deactivate a next hierarchy of application features (435). The computing device 10 may then continue the monitoring accordingly (410).

[0052] It is contemplated that the loop from process steps (435), (410), and (415) can continue until the threat of an out-of-memory crash is alleviated, or until the computing device 100 returns to nominal memory conditions. It is further contemplated that the memory pressure thresholds need not be exceeded sequentially, but rather a prediction can indicate that any memory pressure threshold will be exceeded at a future instance, and can perform the memory reduction actions on all corresponding hierarchically tiered application features, as described herein.

[0053] HARDWARE DIAGRAM

[0054] FIG. 5 is a block diagram that illustrates a computing device upon which examples described herein may be implemented. In various implementations, the computing device 500 can comprise a mobile computing device, such as a smartphone, tablet computer, laptop computer, VR or AR headset device, and the like. As such, the computing device 500 can include telephony features such as a microphone 545, a camera 550, and a communication interface 510 to communicate with external entities using any number of wireless communication protocols. The computing device 500 can further include a positioning module 560 (e.g., a GPS receiver) and an inertial measurement unit 564 that includes one or more accelerometers, gyroscopes, or magnetometers.

[0055] In certain aspects, the computing device 500 can store a service application 532 and an application manager 534 in a local memory 530. In variations, the memory 530 can store additional applications executable by one or more processors 540 of the computing device 500, enabling access and interaction with one or more host servers over one or more networks 580. In certain examples, the application manager 534 can comprise memory monitoring and application managing functionality as described throughout the present disclosure. Furthermore, the application manager 534 may be a component of the service application 532 itself, or integrated as a memory management feature of the service application 532. In such examples, the application manager 534 can receive, as input, memory warnings from an OS low memory warning system of the computing device 500, as well as memory usage metrics of the service application 532 (e.g., historical memory usage data and current memory usage data).

[0056] In variations, the application manager 534 can comprise a component of the operating system of the computing device 500. As such, it is contemplated that the application manager 534 can function to monitor all executing applications on the computing device 500 can implement memory management techniques described herein for each executing application.

[0057] In response to a user input 518 by a user (e.g., the selection of an icon representing the designated application 532), the service application 532 can be executed by a processor 540, which can cause a display interface 522 to be generated on a display screen 520 of the computing device 500. In various implementations, the display interface 522 can enable a user to view and/or interact with displayed content on the display interface 522. Furthermore, the application manager 534 can be initiated based on the service application 532 being launched.

[0058] In various examples, the positioning module 560 can provide location data indicating the current location of the user to the remote computing system 590 (e.g., a host server for the service application 532) to, for example, enable the computing system 590 to provide the display interface 522 with content of the service application 532. In further implementations, the computing device 500 can also transmit additional sensor data to the computing system 590, such as image data from the camera 550, audio data from the microphone 545, and IMU data from the IMU 564 (e.g., continuously or periodically).

[0059] Examples described herein are related to the use of the computing device 500 for implementing the techniques described herein. According to one example, those

techniques are performed by the computing device 500 in response to the processor 540 executing one or more sequences of one or more instructions contained in the main memory 530. Such instructions may be read into the main memory 530 from another machine-readable medium, such as a storage device. Execution of the sequences of instructions contained in the main memory 530 causes the processor 510 to perform the process steps described herein. In alternative implementations, hard-wired circuitry may be used in place of or in combination with software instructions to implement examples described herein. Thus, the examples described are not limited to any specific combination of hardware circuitry and software.

[0060] It is contemplated for examples described herein to extend to individual elements and concepts described herein, independently of other concepts, ideas or systems, as well as for examples to include combinations of elements recited anywhere in this application. Although examples are described in detail herein with reference to the accompanying drawings, it is to be understood that the concepts are not limited to those precise examples. As such, many modifications and variations will be apparent to practitioners skilled in this art. Accordingly, it is intended that the scope of the concepts be defined by the following claims and their equivalents. Furthermore, it is contemplated that a particular feature described either individually or as part of an example can be combined with other individually described features, or parts of other examples, even if the other features and examples make no mention of the particular feature. Thus, the absence of describing combinations should not preclude claiming rights to such combinations.

WHAT IS CLAIMED IS:

1. A computing device comprising:
a display screen;
one or more processors; and
one or more memory resources storing instructions that, when executed by the one or more processors, cause the computing device to:
monitor a set of memory usage metrics of the computing device;
based on historical memory usage data and the set of memory usage metrics, determine whether memory usage will exceed a critical memory threshold at a future instance in time; and
in response to determining that the memory usage will exceed the critical memory threshold at the future instance in time, degrade one or more application features of an application executing on the computing device.
2. The computing device of claim 1, wherein the executed instructions further cause the computing device to:
display streaming content of the application on the display screen;
wherein the executed instructions cause the computing device to degrade the one or more application features by reducing a content render rate of the streaming content.
3. The computing device of claim 1, wherein the application comprises a navigation application, and wherein the executed instructions further cause the computing device to:
display navigation content of the navigation application on the display screen;
wherein the executed instructions cause the computing device to degrade the one or more application features by reducing a resolution of the navigation content.
4. The computing device of claim 1, wherein the set of memory usage metrics comprises a total memory size of the computing device and a current available memory of the computing device.
5. The computing device of claim 1, wherein the executed instructions cause the computing device to collect the historical memory usage data over a predetermined

period prior to determining whether memory usage will exceed the critical memory threshold at the future instance in time.

6. The computing device of claim 5, wherein the executed instructions cause the computing device to determine whether the memory usage will exceed the critical memory threshold at the future instance in time by projecting the historical memory usage data and current memory usage data over a future period of time.

7. The computing device of claim 1, wherein the application comprises multiple priority tiers of application features.

8. The computing device of claim 7, wherein the executed instructions cause the computing device to degrade the one or more application features by degrading a first set of application features in a first priority tier of the multiple priority tiers.

9. The computing device of claim 8, wherein the executed instructions further cause the computing device to:

while or subsequent to degrading the first set of application features, determine that the memory usage at the future instance in time will still exceed the critical memory threshold; and

based on determining that the memory usage at the future instance in time will still exceed the critical memory threshold, degrade a second set of application features in a second priority tier of the multiple priority tiers.

10. A non-transitory computer readable medium storing instructions that, when executed by one or more processors of a computing device, cause the computing device to:

monitor a set of memory usage metrics of the computing device;

based on historical memory usage data and the set of memory usage metrics, determine whether memory usage will exceed a critical memory threshold at a future instance in time; and

in response to determining that the memory usage will exceed the critical memory threshold at the future instance in time, degrade one or more application features of an application executing on the computing device.

11. The non-transitory computer readable medium of claim 10, wherein the executed instructions further cause the computing device to:

display streaming content of the application on the display screen;

wherein the executed instructions cause the computing device to degrade the one or more application features by reducing a content render rate of the streaming content.

12. The non-transitory computer readable medium of claim 10, wherein the application comprises a navigation application, and wherein the executed instructions further cause the computing device to:

display navigation content of the navigation application on the display screen;

wherein the executed instructions cause the computing device to degrade the one or more application features by reducing a resolution of the navigation content.

13. The non-transitory computer readable medium of claim 10, wherein the set of memory usage metrics comprises a total memory size of the computing device and a current available memory of the computing device.

14. The non-transitory computer readable medium of claim 10, wherein the executed instructions cause the computing device to collect the historical memory usage data over a predetermined period prior to determining whether the memory usage will exceed the critical memory threshold at the future instance in time.

15. The non-transitory computer readable medium of claim 14, wherein the executed instructions cause the computing device to determine whether the memory usage will exceed the critical memory threshold at the future instance in time by projecting the historical memory usage and current memory usage over a future period of time.

16. The non-transitory computer readable medium of claim 10, wherein the application comprises multiple priority tiers of application features.

17. A computer-implemented method of managing memory, the method being performed by one or more processors of a computing device and comprising:

monitoring a set of memory usage metrics of the computing device;

based on historical memory usage data and the set of memory usage metrics, determining whether memory usage will exceed a critical memory threshold at a future instance in time; and

in response to determining that the memory usage will exceed the critical memory threshold at the future instance in time, degrading one or more application features of an application executing on the computing device.

18. The method of claim 17, wherein the executed instructions further cause the computing device to:

displaying streaming content of the application on the display screen;

wherein the one or more processors degrade the one or more application features by reducing a content render rate of the streaming content.


19. The method of claim 17, wherein the application comprises a navigation application, the method further comprising:

displaying navigation content of the navigation application on the display screen;

wherein the one or more processors degrade the one or more application features by reducing a resolution of the navigation content.

20. The method of claim 17, wherein the set of memory usage metrics comprise a total memory size of the computing device and a current available memory of the computing device.

FIG. 1



2/5






FIG. 2


3/5

**FIG. 3**

4/5

**FIG. 4**

5/5

**FIG. 5**

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2019/058526

A. CLASSIFICATION OF SUBJECT MATTER

G06F 9/50 (2006.01) G06F 11/34 (2006.01)

According to International Patent Classification (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)

FAMPAT: memory, storage, flash, SSD, RAM, 内存, 存储器, 闪存, 固态硬盘, usage, capacity, 使用量, 容量, historical, past, previous, data, record, 历史, 过去, 以前, 从前, 数据, 记录, exceed, surpass, over, 超过, 超出, threshold, limit, maximum, 临界, 门限, 阈值, 极限, out of memory, memory leak, 内存不足, 内存溢出, monitor, predict, estimate, 监测, 预估, 估计, 预测, degrade, reduce, deactivate, eliminate, remove, stop, cease, suspend, pause, 降低, 减少, 删除, 停止, 停用, 退出, 推迟, 暂停, 关闭, application, program, 应用, 程序, 软件, feature, speed, rate, resolution, process, module, function, 功能, 特征, 速率, 分辨率, 进程, 任务, 项目, 模块, render, stream, display, 成像, 渲染, 数据流 and other related terms.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2016/0036722 A1 (OBRECHT M. E. ET AL.) 4 February 2016 Paragraphs [0009], [0048], [0060], [0066], [0069], [0108], [0170], [0173], [0182]-[0184]	1-20
X	US 10101910 B1 (KARPPANEN J. J.) 16 October 2018 Figures 2, 4; col. 2, lines 38-60; col. 3, lines 5-8; col. 30, lines 45-67	1-20
X	CN 103577306 A (INSPUR) 12 February 2014 Paragraphs [0005]-[0007], [0025] of the original non-English language document (a machine translation is enclosed only for your reference)	1-20

☒ 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

"D" document cited by the applicant in the international application

"E" earlier application or patent but published on or after the international filing date

"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)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"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

"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

"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

"&" document member of the same patent family

Date of the actual completion of the international search

14/01/2020 (day/month/year)

Date of mailing of the international search report

22/01/2020 (day/month/year)

Name and mailing address of the ISA/SG



Intellectual Property Office of Singapore
1 Paya Lebar Link, #11-03
PLQ 1, Paya Lebar Quarter
Singapore 408533

Email: pct@ipos.gov.sg

Authorized officer

Cai Chengjie (Dr)

IPOS Customer Service Tel. No.: (+65) 6339 8616

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2019/058526**C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2018/0129537 A1 (KURTZMAN T. ET AL.) 10 May 2018 Paragraph [0044]	
A	US 2008/0301504 A1 (CHEN J. ET AL.) 4 December 2008 Paragraph [0035]	

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2019/058526

Note: This Annex lists known patent family members relating to the patent documents cited in this International Search Report. This Authority is in no way liable for these particulars which are merely given for the purpose of information.

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2016/0036722 A1	04/02/2016	US 9098333 B1	04/08/2015
US 10101910 B1	16/10/2018	NONE	
CN 103577306 A	12/02/2014	NONE	
US 2018/0129537 A1	10/05/2018	EP 3538995 A1 WO 2018/089317 A1 CN 109923524 A	18/09/2019 17/05/2018 21/06/2019
US 2008/0301504 A1	04/12/2008	NONE	