

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property
Organization
International Bureau



(10) International Publication Number

WO 2018/198002 A1

(43) International Publication Date
01 November 2018 (01.11.2018)

(51) International Patent Classification:

H04N 13/144 (2018.01) *G02B 27/01* (2006.01)
A63F 13/52 (2014.01) *G09G 5/14* (2006.01)

(21) International Application Number:

PCT/IB2018/052735

(22) International Filing Date:

19 April 2018 (19.04.2018)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

15/496,972 25 April 2017 (25.04.2017) US

(71) **Applicant:** ATI TECHNOLOGIES ULC [CA/CA]; One Commerce Valley Dr. East, Markham, Ontario L3T 7X6 (CA).

(72) **Inventor:** RIGUER, Guennadi; One Commerce Valley Dr. East, Markham, L3T 7X6 (CA).

(74) **Agent:** SMITH, Ryan T. et al.; Perry + Currier, 1300 Yonge Street, Suite 500, Toronto, Ontario M4T 1X3 (CA).

(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).

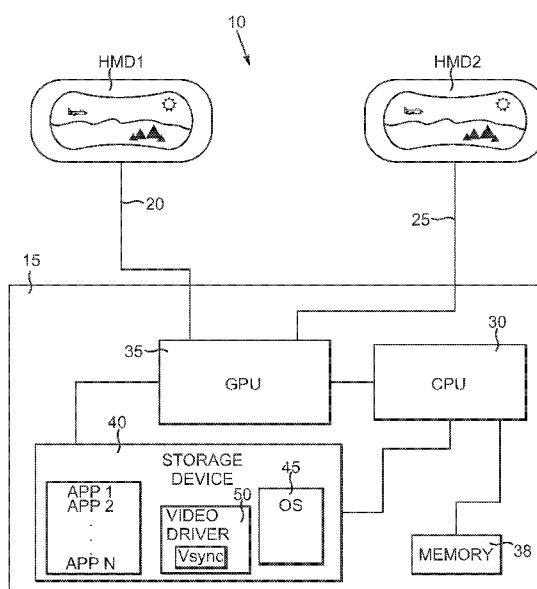
Declarations under Rule 4.17:

- as to the identity of the inventor (Rule 4.17(i))
- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))

Published:

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

(54) Title: DISPLAY PACING IN MULTI-HEAD MOUNTED DISPLAY VIRTUAL REALITY CONFIGURATIONS



(57) **Abstract:** Various virtual reality computing systems and methods are disclosed. In one aspect, a method of delivering video frame data to multiple VR displays (HMD1, HMD2) is provided. The method includes generating content for multiple VR displays and sensing for competing needs for resources with real time requirements of the multiple VR displays. If competing needs for resources with real time requirements are sensed, a selected refresh offset for refreshes of the multiple VR displays is determined to avoid conflict between the competing needs for resources of the multiple VR displays. The selected refresh offset is imposed and the content is delivered to the multiple VR displays.

FIG. 1

DISPLAY PACING IN MULTI-HEAD MOUNTED DISPLAY VIRTUAL REALITY CONFIGURATIONS

CROSS-REFERENCE TO RELATED APPLICATIONS

5 [0001] This application claims priority from U.S. Patent Application No. 15/496,972, filed on April 25, 2017, the entire contents of which are hereby incorporated by reference herein.

BACKGROUND OF THE INVENTION

[0002] Virtual Reality (VR) systems utilize a Head Mounted Display (HMD) with an illuminated screen that shows the wearer the computer generated virtual world. With each movement of a user's head the HMD changes the display to show the user the changed scenery. When running an application, such as a game, a VR computer must not only render video frames generated by the application but also account for the movements of a user's head. With or without head movement, the VR computer must process frames in near real time in order to avoid judder, which results when images are delayed and the user will view the images just a second, or even a fraction of a second too late. Judder can be very jarring and cause nausea.

BRIEF DESCRIPTION OF THE DRAWINGS

[0003] The foregoing and other advantages of the invention will become apparent upon reading the following detailed description and upon reference to the drawings in which:

[0004] FIG. 1 is a schematic view of an exemplary VR computing system;

5 [0005] FIG. 2 depicts a conventional exemplary timing diagram for multiple VR display refresh cycles;

[0006] FIG. 3 depicts an exemplary timing diagram for multiple VR display refresh cycles for an exemplary variant;

[0007] FIG. 4 depicts an exemplary timing diagram for multiple VR display refresh cycles for another exemplary variant;

10 [0008] FIG. 5 is a schematic view like FIG. 1, but depicting an alternate exemplary VR computing system;

[0009] FIG. 6 is a flow chart depicting an exemplary process flow for delivering frames to one or more VR displays using selected refresh offset or not; and

15 [0010] FIG. 7 is a flow chart depicting an exemplary process flow for delivering frames to multiple VR displays using selected refresh offset and dynamic refresh.

DETAILED DESCRIPTION

[0011] A VR system relies on an ability to perform in near real-time to provide a good user experience. As noted above, a single HMD driven system represents a straightforward situation. Complexities mount when multiple HMDs are connected to a computer system. A conventional VR HMD operates on a fixed time line usually tied to a display refresh rate. Vertical synchronization or Vsync is a method of controlling buffer flipping to avoid the problem of frame tearing where the display shows portions of two rendered frames that have different content. When Vsync is enabled, each buffer flip can only occur after each refresh. This effectively caps the frame rendering rate at the HMD refresh rate.

[0012] Some VR system vendors introduced a technique called asynchronous time warp (ATW) to help alleviate the impact of missed frames on VR HMDs. A system using ATW shifts the image on the VR HMD without input from the game engine based on relative head motion that occurred after the last VR HMD position was sent to the game. Doing so presents a more accurate image to the user. ATW moves the entire frame as a whole, shifting it only based on relative changes to the user's head rotation. A more recent development is the so-called Asynchronous Spacewarp (ASW). Systems with ASW attempt to shift objects and motion inside of the scene by generating new frames to insert in between "real" frames from the game engine when the game is running at a frame rate lower than the refresh rate.

[0013] ATW is executed at a minimal latency to the front portion of the Vsync interval. Similarly, other parts of the frame generation executed on both the central processing unit (CPU) and the graphics processing unit (GPU) are paced relative to HMD's Vsync. This works well for a system with a single HMD. When multiple HMDs are attached, even if the system (e.g. CPU and GPU) has enough total horsepower, unreliable operation could occur due to contention for real-time critical resources (e.g. computations for ATW) across multiple HMDs. There are conventional solutions to synchronize display processing and output, such as gen lock and frame lock, but there are no solutions for pacing display operation in a predictable manner as required for multiple VR headsets.

[0014] In accordance with one aspect of the present invention, a method of delivering video frame data to multiple VR displays is provided. The method includes generating content for multiple VR displays and sensing for competing needs for resources with real time requirements of the multiple VR displays. If competing needs for resources with real time requirements are sensed, a selected refresh offset for refreshes of the multiple VR displays is determined to avoid conflict between the competing needs for resources of the multiple VR displays. The selected refresh offset is imposed and the content is delivered to the multiple VR displays.

[0015] The method wherein the resources comprise computation for rendering and asynchronous time warp requests.

[0016] The method wherein the multiple displays support dynamic refresh and comprising if competing needs for resources with real time requirements are sensed, also determining a selected dynamic refresh rate for refreshes of the multiple VR displays to aid in avoiding the competing needs for resources made by the multiple VR, and imposing the selected refresh offset and dynamic refresh rate.

5 [0017] The method wherein the resources comprise computation for rendering and asynchronous time warp requests.

[0018] The method wherein the generating the content is performed by a single GPU.

10 [0019] The method of claim wherein the generating the content for multiple VR displays comprises generating the content for one of the multiple VR displays using a GPU and generating or delivering the content for another of the multiple VR displays using another GPU.

[0020] The method wherein the GPU is configured as a master and the another GPU is configured as a slave such that the master controls the selected refresh offset of frames generated or delivered by the slave GPU.

15 [0021] In accordance with another aspect of the present invention, a method of delivering video frame data to multiple VR displays is provided. The method includes running a first application on a computing device to generate content for multiple VR displays and sensing for competing needs for resources with real time requirements of the multiple VR displays using a second application. If competing needs for resources with real time requirements are sensed, using the second application to determine a selected refresh offset for refreshes of the multiple VR displays to avoid conflict between competing needs for resources of the multiple VR displays. The selected refresh offset is imposed and the content delivered to the multiple VR displays.

[0022] The method wherein the resources comprise computation for rendering and asynchronous time warp requests.

25 [0023] The method wherein the multiple displays support dynamic refresh, the method comprising if movements are sensed, also determining a selected dynamic refresh rate for refreshes of the multiple VR displays to aid in avoiding the competing requests for resources made by the multiple VR displays due to the movements, and imposing the selected refresh offset and dynamic refresh rate.

[0024] The method wherein the resources comprises computation for rendering and asynchronous time warp requests.

30 [0025] The method wherein the application is run by a single GPU.

[0026] The method wherein the generating the content for multiple VR displays comprises generating the content for one of the multiple VR displays using a GPU and generating or delivering the content for another of the multiple VR displays using another GPU.

[0027] The method wherein the GPU is configured as a master and the another GPU is configured as a slave such that the master controls the selected refresh offset of frames generated or delivered by the slave GPU.

[0028] In accordance with another aspect of the present invention, a virtual reality computing system is provided. The system includes a computing device and a processor. The processor is operable to perform instructions to generate content for multiple VR displays, to sense for competing needs for resources with real time requirements of the multiple VR displays, if competing needs for resources with real time requirements are sensed, to determine a selected refresh offset for refreshes of the multiple VR displays to avoid conflict between the competing needs for resources of the multiple VR displays, to impose the selected refresh offset, and to deliver the content to the multiple VR displays.

[0029] The virtual reality computing system comprising the multiple VR displays.

[0030] The virtual reality computing system wherein the processor comprises a CPU, a GPU or a combined CPU and GPU.

[0031] The virtual reality computing system wherein the computing device comprises another processor wherein the processor generates the content for one of the multiple VR displays and the another processor generates or delivers the content for another of the multiple VR displays.

[0032] The virtual reality computing system wherein the processor is configured as a master and the another processor is configured as a slave such that the master controls the selected refresh offset of frames generated or delivered by the slave processor.

[0033] The virtual reality computing system wherein the multiple displays support dynamic refresh, the processor being operable to, if competing needs for resources with real time requirements are sensed, also determine a selected dynamic refresh rate for refreshes of the multiple VR displays to aid in avoiding conflict between the competing needs for resources of the multiple VR displays, and to impose the selected refresh offset and dynamic refresh rate.

[0034] Various virtual reality (VR) computing systems that improve on the management of delivering video frames to multiple VR head mounted displays are disclosed. Video content can be generated for delivery to multiple VR head mounted displays. The VR computer monitors, perhaps by way of a VR compositor or other application for competing tasks requiring real time processing, such as partial or whole frame rendering or asynchronous time warps, by multiple VR displays. If competing tasks that require real time or near real time computing requirements are sensed, say by the detection of multiple VR displays, the application and the system compute a selected refresh offset for refreshes of the multiple VR displays to avoid competing requests for resources made by the multiple VR displays. The selected refresh offset is imposed and the content is delivered to the multiple VR displays. One head mounted VR display can refresh first, and the second a fraction of a second or a fraction of a frame later and so on for

other VR displays connected to the computing device. The staggering avoids competing requests for resources, such as those needed for asynchronous time warps and frame rendering parts that have real time requirements. Additional details will now be described.

[0035] In the drawings described below, reference numerals are generally repeated where identical elements appear in more than one figure. Turning now to the drawings, and in particular to FIG. 1, therein is shown a schematic view of an exemplary VR computing system 10 (hereinafter system 10). The system 10 includes a computing device 15 and one or more head mounted displays, two of which are depicted and labeled HMD1 and HMD2, respectively, where HMD stands for head mounted display. The connections 20 and 25 between the computing device 15 and HMD1 and HMD2 can be wired or wireless connections as desired. The computing device 15 can include a processor or CPU 30, a GPU 35, system memory 38 and a storage device 40. The CPU 30 and the GPU 35 can be general purpose discrete devices or integrated into a single device, such as by way of an accelerated processing unit. Optionally, the functions of the CPU 30 and the GPU 35 can be performed by one or more application specific integrated circuits. System memory can be RAM, ROM, flash combinations of these or other types of memory. The storage device 40 is a non-volatile computer readable medium and can be any kind of hard disk, optical storage disk, solid state storage device, ROM, RAM or virtually any other system for storing computer readable media. The storage device 40 is operable to store non-transient computer readable instructions for performing various functions disclosed herein.

[0036] The computing device 15 can include plural applications, which are abbreviated APP 1, APP 2 ... APP N, and which can be drivers, software applications, or other types of applications. In addition, the computing device 15 can include an operating system 45 and a video driver 50. The operating system 45 and the video driver 50 and the applications APP 1, APP 2 ... APP n can be stored on the storage device 40 and selectively loaded into system memory 38. Windows®, Linux, or more application specific types of operating system software can be used or the like. One of the applications APP 1 ... APP n, for example APP 1, can be an application that generates content to be displayed on HMD1 and HMD2, such as a game or other program, and another application, say APP 2, can be a VR compositor that, by way of the video driver 50, controls aspects of how content is displayed on HMD1 and HMD2. The video driver 50 can include code to implement Vsync or vertical synchronization. Vsync is designed to avoid a phenomena known as “tearing” where the display shows portions of two rendered frames that have different content. When Vsync is enabled, each buffer flip can only occur after each refresh.

[0037] The application APP 1 and/or the application APP 2 can be resident on the storage device 40 and selectively loaded in the system memory 38 during execution. The application APP 2 is operable to provide timing synchronization for driving the displays of HMD1 and HMD2 as well as selectively skewing the refresh intervals of the multiple VR displays for HMD1 and HMD2. A technical goal is to

provide some selected offset in refresh timing between HMD1 and HMD2 to avoid overlap of real time workloads, such as frame rendering and asynchronous time warp (ATW) requests. Aspects of the operation of the system 10 will be described in conjunction with FIG. 3. However, before turning to FIG. 3, it will be useful to consider conventional frame rendering in a multi-HMD system. Such conventional frame rendering is depicted in FIG. 2, which illustrates a refresh timing diagram 55 for HMD1 and a corresponding refresh timing diagram 60 for HMD2. For each of the timing diagrams 55 and 60 a few refresh points, Refresh 1, Refresh 2, Refresh 3 and Refresh 4, are depicted and it is assumed that both HMD1 and HMD2 have the same constant refresh rate, i.e., the inverse of the time period between Refresh 1 and Refresh 2 and between Refresh 2 and Refresh 3 and so on. The refresh timing diagram 55 shows timing activities for a “Hardware Queue” and a “Displayed On HMD1” where the Hardware Queue is the hardware (e.g., GPU, APU or other) activity associated with frame rendering for HMD1 and the Displayed On HMD1 shows the activity leading to the actual appearance of rendered content on HMD1. As shown in diagram 55, at some point after Refresh 1 but prior to Refresh 2, frame rendering for HMD1, abbreviated “Render 1,” is performed and the content rendered just prior to Refresh 1 is displayed on HMD1, abbreviated “Send to HMD1.” The refresh timing shows the timing activities for the Hardware Queue associated with frame rendering for HMD2 and the “Displayed On HMD2” shows the activity leading to the actual appearance of rendered content on HMD2. As shown in diagram 60, at some point after Refresh 1 but prior to Refresh 2, frame rendering for HMD2, abbreviated “Render 2,” is performed and the content rendered just prior to Refresh 1 is displayed on HMD2, abbreviated “Send to HMD2.” Here it is assumed that Render 2 commences at time t_0 very soon after Refresh 1 and that Render 1 commences at time t_1 immediately and sequentially after the completion of Render 2. Now fast forward to just after Refresh 2. As shown in refresh timing diagram 55, the content from the Render 1 operation commencing at time t_1 is ultimately displayed on HMD1 at the completion of Send to HMD1 at time t_2 . Note the latency, labeled “Latency HMD1,” between Render 1 commencing at time t_1 and Send to HMD1 at time t_2 . Now consider timing diagram 60 just after Refresh 2. The content from the Render 2 operation commencing at time t_0 is ultimately displayed on HMD2 at the completion of Send to HMD2 at time t_2 . But note the latency, labeled “Latency HMD2,” between Render 2 commencing at time t_0 and Send to HMD2 at time t_2 . Render 2 and Render 1 compete for processing between Refresh 1 and Refresh 2 and thus must be staggered in time. This means that Render 2 must be performed early in the Refresh 1 to Refresh 2 period and this results in Latency HMD2>Latency HMD1 and perhaps Latency HMD2 being long enough to cause judder for the user of HMD2.

[0038] Attention is now turned to FIG. 3, which illustrates exemplary improved operations and in particular a refresh timing diagram 65 for HMD1 and a corresponding refresh timing diagram 70 for HMD2. For each of the timing diagrams 65 and 70 a few refresh points, Refresh 1, Refresh 2, Refresh 3

and Refresh 4, are depicted. The refresh rate R_{HMD1} of HMD1 is given by $1/T_{HMD1}$ where T_{HMD1} is the period between Refresh 1 and Refresh 2 and between Refresh 2 and Refresh 3 and so on in the diagram 65 for HMD1. The refresh rate R_{HMD2} of HMD2 is given by $1/T_{HMD2}$ where T_{HMD2} is the period between Refresh 1 and Refresh 2 and between Refresh 2 and Refresh 3 and so on in the diagram 60 for HMD2.

5 The refresh rates R_{HMD1} and R_{HMD2} can be the same or different and can be dynamically adjusted if HMD1 and HMD2 and the computing device 15 are capable of dynamic refresh rates. The refresh timing diagram 65 shows timing activities for a “Hardware Queue” and a “Displayed On HMD1” where the

Hardware Queue is the hardware (e.g., GPU, APU or other) activity associated with frame rendering for HMD1 and the Displayed On HMD1 shows the activity leading to the actual appearance of rendered content on HMD1. As shown in diagram 65, at time t_4 after Refresh 1 but prior to Refresh 2, frame rendering for HMD1, abbreviated “Render 1,” is commenced and the content rendered just prior to Refresh 1 is displayed on HMD1, abbreviated “Send to HMD1.” Note that time t_4 is selected so that Render 1 completes near Refresh 2. Now fast forward to just after Refresh 2 in diagram 65. As shown in refresh timing diagram 65, the content from the Render 1 operation commencing at time t_4 is ultimately displayed on HMD1 at the completion of Send to HMD1 at time t_5 . Note the latency, labeled “Latency HMD1,”

10 between Render 1 commencing at time t_4 and Send to HMD1 at time t_5 . Now consider timing the refresh timing diagram 70. The refresh timing diagram 70 shows timing activities for the Hardware Queue and a “Displayed On HMD2” where the Hardware Queue is the hardware (e.g., GPU, APU or other) activity associated with frame rendering for HMD2 and the Displayed On HMD2 shows the

15 activity leading to the actual appearance of rendered content on HMD2. But note that Refresh 1 for the Hardware Queue in refresh timing diagram 70, in other words the Refresh 1 for HMD2, is offset in time relative to the Refresh 1 for HMD1 shown in refresh timing diagram 65 by some selected Refresh Offset.

20 As shown in diagram 70, at time t_6 after Refresh 1 but prior to Refresh 2, frame rendering for HMD2, abbreviated “Render 2,” is commenced and the content rendered just prior to Refresh 1 is displayed on HMD2, abbreviated “Send to HMD2.” Note that time t_6 is selected so that Render 2 completes near

25 Refresh 2. Now fast forward to just after Refresh 2 in diagram 70. As shown in refresh timing diagram 70, the content from the Render 2 operation commencing at time t_6 is ultimately displayed on HMD2 at the completion of Send to HMD2 at time t_7 . Note the latency, labeled “Latency HMD2,” between

30 Render 1 commencing at time t_6 and Send to HMD2 at time t_7 . By using the Refresh Offset, Render 1 can be completed near Refresh 2 for HMD1 and Render 2 can be completed near Refresh 2 for HMD2, which results in Latency HMD1 being equal or close to Latency HMD2 and Latency HMD2 being shorter than Latency HMD2 for the conventional technique illustrated in FIG. 2.

[0039] Still referring to FIG. 3, the Refresh Offset can be selected based on some fraction (say up to one half) of a frame up or some other software controlled interval such as 1 to 2 milliseconds, etc. The purpose of the selected Refresh Offset is to space out in time HMD1 and HMD2 tasks with real time

requirements, such as ATW requests and/or partial or whole frame rendering. It should be understood that the refresh rates R_{HMD1} and R_{HMD2} can not only be different but also fluctuate rapidly if dynamic refresh rates are supported by the HMD1 and HMD2 and the computing device 15.

[0040] Some exemplary selections of Refresh Offset will now be described in conjunction with FIGS. 1 and 3. One exemplary technique involves a brute force selection of the maximum Refresh Offset. For example, the maximum Refresh Offset in terms of a fraction of a frame will be given by the inverse of the number of head mounted displays connected to the computing device 15. So, for example, where HMD1 and HMD2 are connected to the computing device 15, the maximum Refresh Offset in terms of frames will be $\frac{1}{2}$ of a frame. If there were three HMDs the maximum Refresh Offset would be $\frac{1}{3}$ of a frame and for four HMDs the maximum Refresh Offset would be $\frac{1}{4}$ and so on. The selection and imposition of the Refresh Offset(s) can be performed by the VR compositor, such as APP 2 in these examples. The Refresh Offset can be adjusted up or down when the computing device 15, via APP 2 or otherwise, senses the connection or disconnection of HMDs.

[0041] Another exemplary technique for selecting the Refresh Offset involves dynamic selection. Here, as the content generating application APP 1 is producing frames for HMD1 and HMD2, the VR compositor (APP 2) can determine how long an exclusive operation will take, and then instruct the video driver 50 what the Refresh Offset should be to permit that exclusive operation to be performed without any competing real time tasks, such as ATW requests or frame rendering overlaps. Immediately after a refresh, say Refresh 1 of HMD1, the VR compositor (APP 2) will poll the tracking data of and thereby sense the movements of HMD1 and HMD2 and also sense any other incoming tasks with real time requirements. With that data in hand, APP 2 will next determine how long the ATW request and frame rendering with the ATW correction for the HMD1 movement will take, and then instruct the video driver 50 what the Refresh Offset should be so that the rendering for the movement of HMD1 is not fettered by ATW requests associated with the movement of the competing HMD, such as HMD2. The APP 2 could instruct the video driver 50 to include a safety margin, that is, make the Refresh Offset slightly larger than necessary. In the next instant, other movements of HMD1 and HMD2 can occur which might prompt APP 2 to readjust the Refresh Offset accordingly. This process of dynamically adjusting the Refresh Offset can occur repeatedly depending on the movements of HMD1 and HMD2. Of course, the technique can be dovetailed with dynamic refresh for all HMDs, such as HMD1 and HMD2, using, for example, Freesync, a technology supplied by Advanced Micro Devices, Inc.

[0042] Certain VR set ups can involve HMDs with different refresh rates. In these circumstances, the APP 2 can still impose Refresh Offsets to space out real time computing tasks. Attention is now turned to FIG. 4, which depicts a refresh timing diagram 75 for HMD1 and a corresponding refresh timing diagram 80 for HMD2. FIG. 4 is like FIG. 3 and thus depicts a few refresh points, Refresh 1, Refresh 2, Refresh 3 and Refresh 4, a Hardware Queue and Displayed on HMD1 or HMD2 traces for each of the

timing diagrams 75 and 80. However, note that the refresh rates, defined as described above, for HMD1 and HMD2 differ. In the absence of any compensations, HMD1 and HMD2 with different refresh rates will produce an irregular “beat” pattern where intervals between refreshes of the multiple HMDs, HMD1 and HMD2, become bigger and smaller and eventually collide. Here, however, APP 2 (as VR

5 compositor, see FIG. 1) can impose a Refresh Offset to attempt to space Refresh 1 for HMD1 and Refresh 1 for HMD2, Refresh 2 for HMD1 and Refresh 2 for HMD2 and so on, such that they do not collide. Other methods of refresh offsetting could also be used. In cases when it is not possible to avoid refresh collisions with multiple HMDs, HMD1 and HMD2, or when there is no collision, but refresh markers from different HMDs are too close to each other, a feedback can be generated to warn APP 2 that it needs to account for a possible increase in latency as a result of a possible collision of workloads.

10 [0043] Multi-GPU variants are envisioned for usage with multiple VR HMDs. In this regard, attention is now turned to FIG. 5 which is a block diagram depicting an alternate exemplary variant of the computing system 10'. This computing system variant 10' shares many of the attributes of the computing system 10 depicted in FIG. 1 and described elsewhere herein. For instance, there can be 15 multiple VR displays HMD1 and HMD2 and these can number more than two as desired. In addition, the computing device 15' can include the GPU 35, system memory 38, the storage device 40, the OS 45 with Vsync, plural applications APP 1 to APP N and a video driver 50. In addition, HMD1 and HMD2 can connect to the computing device 15' by way of the aforementioned connections 20 and 25. Here however, in addition to the GPU 35, the computing device 15' can include one or more additional GPUs, 20 one of which is depicted and labeled 85. It should be understood that more than two GPUs 35 and 85 can be implemented as desired. Both the GPUs 35 and 85 can have operative connections to the storage device 40 and system memory 38 as well as to each other. APP 2, if configured as a VR compositor as 25 described elsewhere herein, can include instructions to perform rendering for HMD1 and HMD2, respectively. For instance, the GPU 35 can be tasked with rendering content for HMD1 and the GPU 85 can be tasked with rendering content for the other VR display HMD2. It can be desirable to designate one of the GPUs, for example GPU 35, as a master and the other GPU or GPUs 85 as slaves for timing 30 synchronization and skewing of refresh rates and points. This differentiation in tasking, that is, master and slave can be performed where the GPU 35 and the GPU 85 have the same capabilities or more commonly where one GPU, say the GPU 35, is a more capable integrated circuit such as a discrete GPU on a card and the other GPU 85 is a less capable but still functional GPU, say for example in an integrated graphics chip. The same type of timing synchronization and selective skewing of refresh intervals as depicted in FIG. 3 and described elsewhere herein can be used with the multiple GPU variant depicted in FIG. 5. Here, however, the master GPU 35 can be tasked with selectively skewing the refresh intervals for the slave GPU 85 and any other slave GPUs that can be incorporated into the computing 35 device 15. It can also be possible to use additional levels of granularity. Other differentiation in tasking

is possible. For example, the more capable GPU 35 can be used for tasks that have real time requirements for both HMD1 and HMD2 and the lesser capable GPU 85 can be used for less demanding tasks, such as driving HMD2. In another variant, the lesser GPU 85 could also be used for some rendering as well. As with the other disclosed variants, dynamic refresh can be implemented using, for 5 example, Freesync or other technologies where dynamic refresh of the displays HMD1 and HMD2 are provided.

[0044] A couple of exemplary process flows are illustrated graphically in FIGS. 6 and 7. FIG. 6 depicts an exemplary process flow without dynamic refresh enabled and FIG. 7 to be discussed below depicts an exemplary process flow with dynamic refresh enabled. FIG. 6 depicts two parallel process 10 loops 100 and 102 that can be asynchronous. Process loop 100 handles the sensing for competing resource requests, and selection and imposition of selected refresh offsets. Thus at step 105, an HMD is added or removed from the computing device 15 shown in FIG. 1. At step 110, an application, for example APP 2 implemented as a VR compositor or similar functionality, senses for competing needs for real time resources of multiple VR displays. This entails sensing the number of VR displays, such as 15 HMD1, HMD2, etc., connected to the computing device 15. A single detected VR display will not generate competing needs for real time resources, but two or more will. If, at step 115, competing needs for real time resources of one or more VR displays are sensed then the process flow proceeds to step 120 and the APP 2 determines an appropriate refresh offset (or offsets) in order to avoid conflict in the needs for real time resources. This determination will take into account how many competing VR displays, 20 such as HMD1, HMD2, etc. are detected--added VR displays will require more refresh offsets and removed VR displays will require fewer refresh offsets. Next at step 125 the APP 2 imposes the refresh offset (or offsets) to avoid conflict in needs for real time resources and then the process proceeds to loop exit at step 130. If, however, at step 115, competing needs for real time resources of one or more VR displays are not sensed then the process flow skips to exit at step 130. The loop is triggered on whenever 25 VR displays are connected or disconnected. The content generation loop 102 proceeds in parallel. Thus following start at step 135, at step 140 the APP 1 generates content for multiple VR displays, such as HMD1, HMD2, etc. Step 145 is a conditional that uses the outcomes of steps 120 and 125 of process loop 100. If at step 145, a Refresh Offset(s) is not imposed, such as would be the case if only one VR display is detected, the loop 102 proceeds to step 150 to render frames without refresh offset, then to step 30 155 to deliver the rendered frames to the single VR display and then back to step 140. If, however, at step 145 a Refresh Offset(s) is imposed, such as would be the case if multiple VR displays are detected, then the loop skips to step 160 and the frames are rendered with refresh offset (or offsets) followed by step 150 where the rendered frames are delivered to the multiple VR displays, such as HMD1, HMD2, etc. The loop 102 then loops back to step 140.

[0045] In an alternate exemplary variant, a process flow can use dynamic refresh by way of Freesync or some other type of dynamic refresh program. Referring now to FIG. 5, after a start at step 200, at step 205 dynamic refresh is enabled. At step 210, an application such as APP 1, generates content for multiple VR displays such as the display HMD1 and HMD2. Next, at step 215, the application APP 2 5 senses for competing needs for real time resources of multiple VR displays. For example, and as described above, APP 2 will sense for competing needs of ATW requests and partial or whole frame rendering of for example, HMD1 and HMD2. If at step 220, no competing needs are detected, then the process flow proceeds to step 225 and the frames are rendered, and then at step 230 the rendered frames are delivered to multiple VR displays. Step 230 is followed by a return to step 210. If, however, at step 10 220, application APP 2 senses competing needs for real time resources of multiple VR displays of one or more of VR displays, then the process proceeds to step 235 where the APP 2 determines a refresh offset and a dynamic refresh rate to avoid conflict in the competing needs for real time resources. Next at step 230 APP 2 imposes the selected refresh offset and dynamic refresh rate to avoid conflict in needs for real 15 time resources, and thereafter at step 245 the frames are rendered with the selected refresh offset and dynamic refresh rate. Next, the process flow branches over to step 230 and the rendered frames are delivered to the multiple VR displays. Of course the process flows in FIGS. 6 and 7 can be accomplished with one or multiple applications and operating systems, etc. In other words, various levels of software integration, and hardware, can be used to implement the described steps.

[0046] While the invention can be susceptible to various modifications and alternative forms, specific 20 variants have been shown by way of example in the drawings and have been described in detail herein. However, it should be understood that the invention is not intended to be limited to the particular forms disclosed. Rather, the invention is to cover all modifications, equivalents and alternatives falling within the spirit and scope of the invention as defined by the following appended claims.

CLAIMS

What is claimed is:

1. A method of delivering video frame data to multiple VR displays (HMD1, HMD2), comprising:
 - 5 generating content for multiple VR displays;
 - sensing for competing needs for resources with real time requirements of the multiple VR displays;
 - 10 if competing needs for resources with real time requirements are sensed, determining a selected refresh offset for refreshes of the multiple VR displays to avoid conflict between the competing needs for resources of the multiple VR displays;
 - imposing the selected refresh offset; and
 - 15 delivering the content to the multiple VR displays.
2. The method of claim 1, wherein the resources comprise computation for rendering and asynchronous time warp requests.
- 15
3. The method of claim 1, wherein the multiple displays support dynamic refresh, the method comprising if competing needs for resources with real time requirements are sensed, also determining a selected dynamic refresh rate for refreshes of the multiple VR displays to aid in avoiding the competing needs for resources made by the multiple VR, and imposing the selected refresh offset and dynamic refresh rate.
- 20
4. The method of claim 3, wherein the resources comprise computation for rendering and asynchronous time warp requests.
- 25
5. The method of claim 1, wherein the generating the content is performed by a single GPU.
6. The method of claim 1, wherein the generating the content for multiple VR displays comprises generating the content for one of the multiple VR displays using a GPU and generating or delivering the content for another of the multiple VR displays using another GPU.
- 30
7. The method of claim 6, wherein the GPU is configured as a master and the another GPU is configured as a slave such that the master controls the selected refresh offset of frames generated or delivered by the slave GPU.
- 35
8. A method of delivering video frame data to multiple VR displays (HMD1, HMD2), comprising:

running a first application (APP 1) on a computing device (15) to generate content for multiple VR displays;

sensing for competing needs for resources with real time requirements using a second application (APP 2);

5 if competing needs for resources with real time requirements are sensed, using the second application to determine a selected refresh offset for refreshes of the multiple VR displays to avoid conflict between the competing needs for resources of the multiple VR displays;

imposing the selected refresh offset; and

10 delivering the content to the multiple VR displays.

9. The method of claim 8, wherein the resources comprise computation for rendering and asynchronous time warp requests.

15 10. The method of claim 8, wherein the multiple displays support dynamic refresh, the method comprising if movements are sensed, also determining a selected dynamic refresh rate for refreshes of the multiple VR displays to aid in avoiding the competing requests for resources made by the multiple VR displays due to the movements, and imposing the selected refresh offset and dynamic refresh rate.

20 11. The method of claim 10, wherein the resources comprises computation for rendering and asynchronous time warp requests.

12. The method of claim 1, wherein the application is run by a single GPU (35).

25 13. The method of claim 8, wherein the generating the content for multiple VR displays comprises generating the content for one of the multiple VR displays using a GPU (35) and generating or delivering the content for another of the multiple VR displays using another GPU (85).

30 14. The method of claim 13, wherein the GPU is configured as a master and the another GPU is configured as a slave such that the master controls the selected refresh offset of frames generated or delivered by the slave GPU.

15. A virtual reality computing system (10), comprising:
35 a computing device (15);

5 a processor (30, 35) operable to perform instructions to generate content for multiple VR displays (HMD1, HMD2), to sense for competing needs for resources with real time requirements of the multiple VR displays, if competing needs for resources with real time requirements are sensed, to determine a selected refresh offset for refreshes of the multiple VR displays to avoid conflict between the competing requests for resources of the multiple VR displays, to impose the selected refresh offset, and to deliver the content to the multiple VR displays.

10 16. The virtual reality computing system of claim 15, comprising the multiple VR displays.

15 17. The virtual reality computing system of claim 15, wherein the processor comprises a CPU, a GPU or a combined CPU and GPU.

20 18. The virtual reality computing system of claim 15, wherein the computing device comprises another processor (85) wherein the processor generates the content for one of the multiple VR displays and the another processor generates or delivers the content for another of the multiple VR displays.

25 19. The virtual reality computing system of claim 18, wherein the processor is configured as a master and the another processor is configured as a slave such that the master controls the selected refresh offset of frames generated or delivered by the slave processor.

20. The virtual reality computing system of claim 15, wherein the multiple displays support dynamic refresh, the processor being operable to, if competing needs for resources with real time requirements are sensed, also determine a selected dynamic refresh rate for refreshes of the multiple VR displays to aid in avoiding conflict between the competing needs for resources of the multiple VR displays, and to impose the selected refresh offset and dynamic refresh rate.

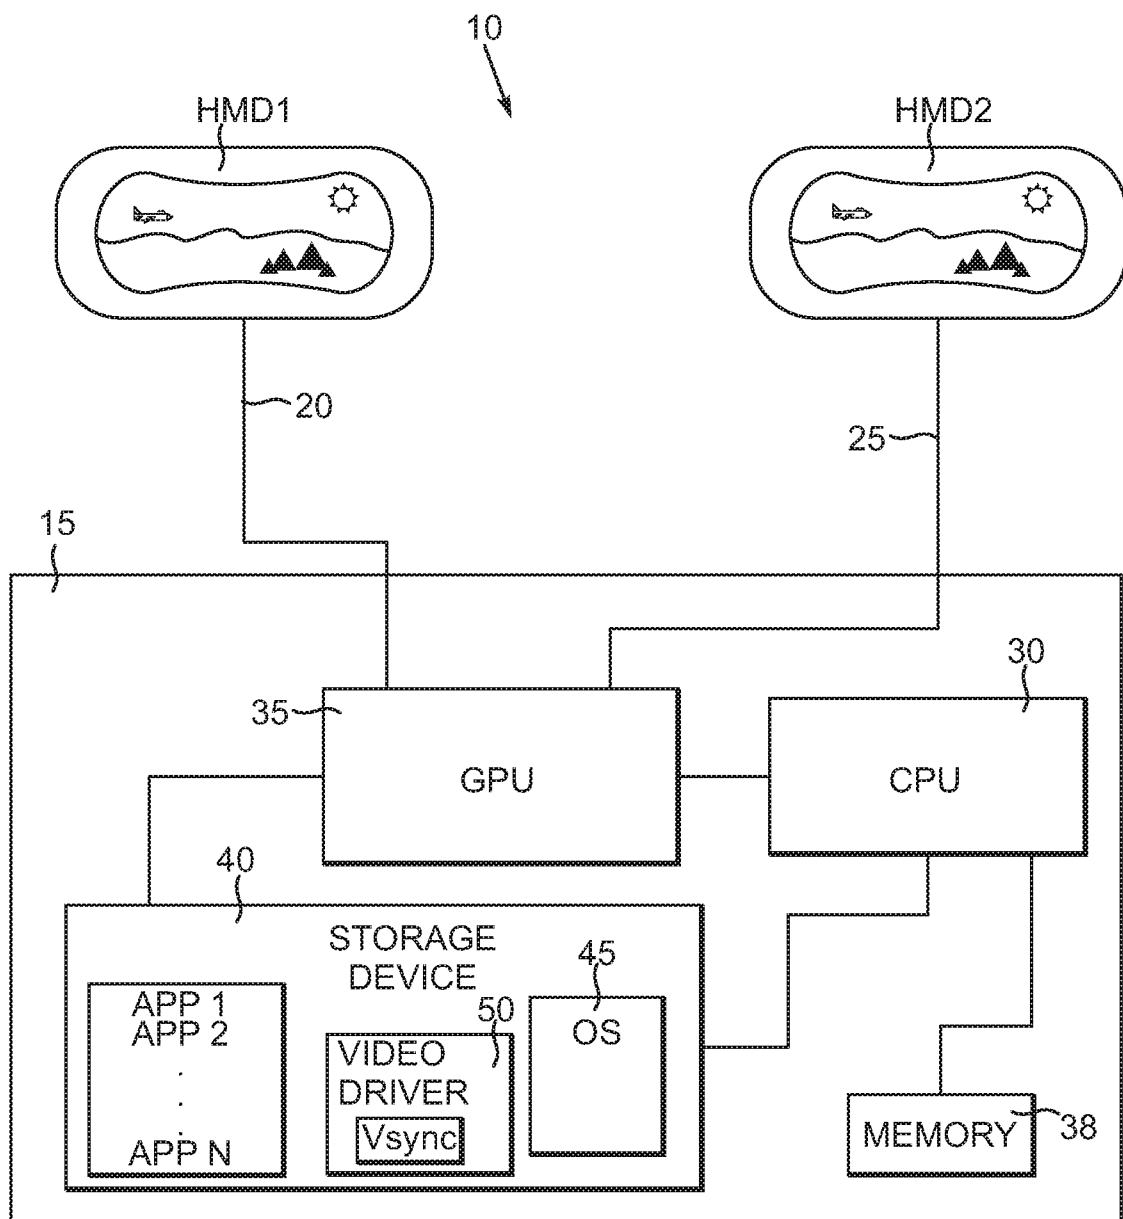


FIG. 1

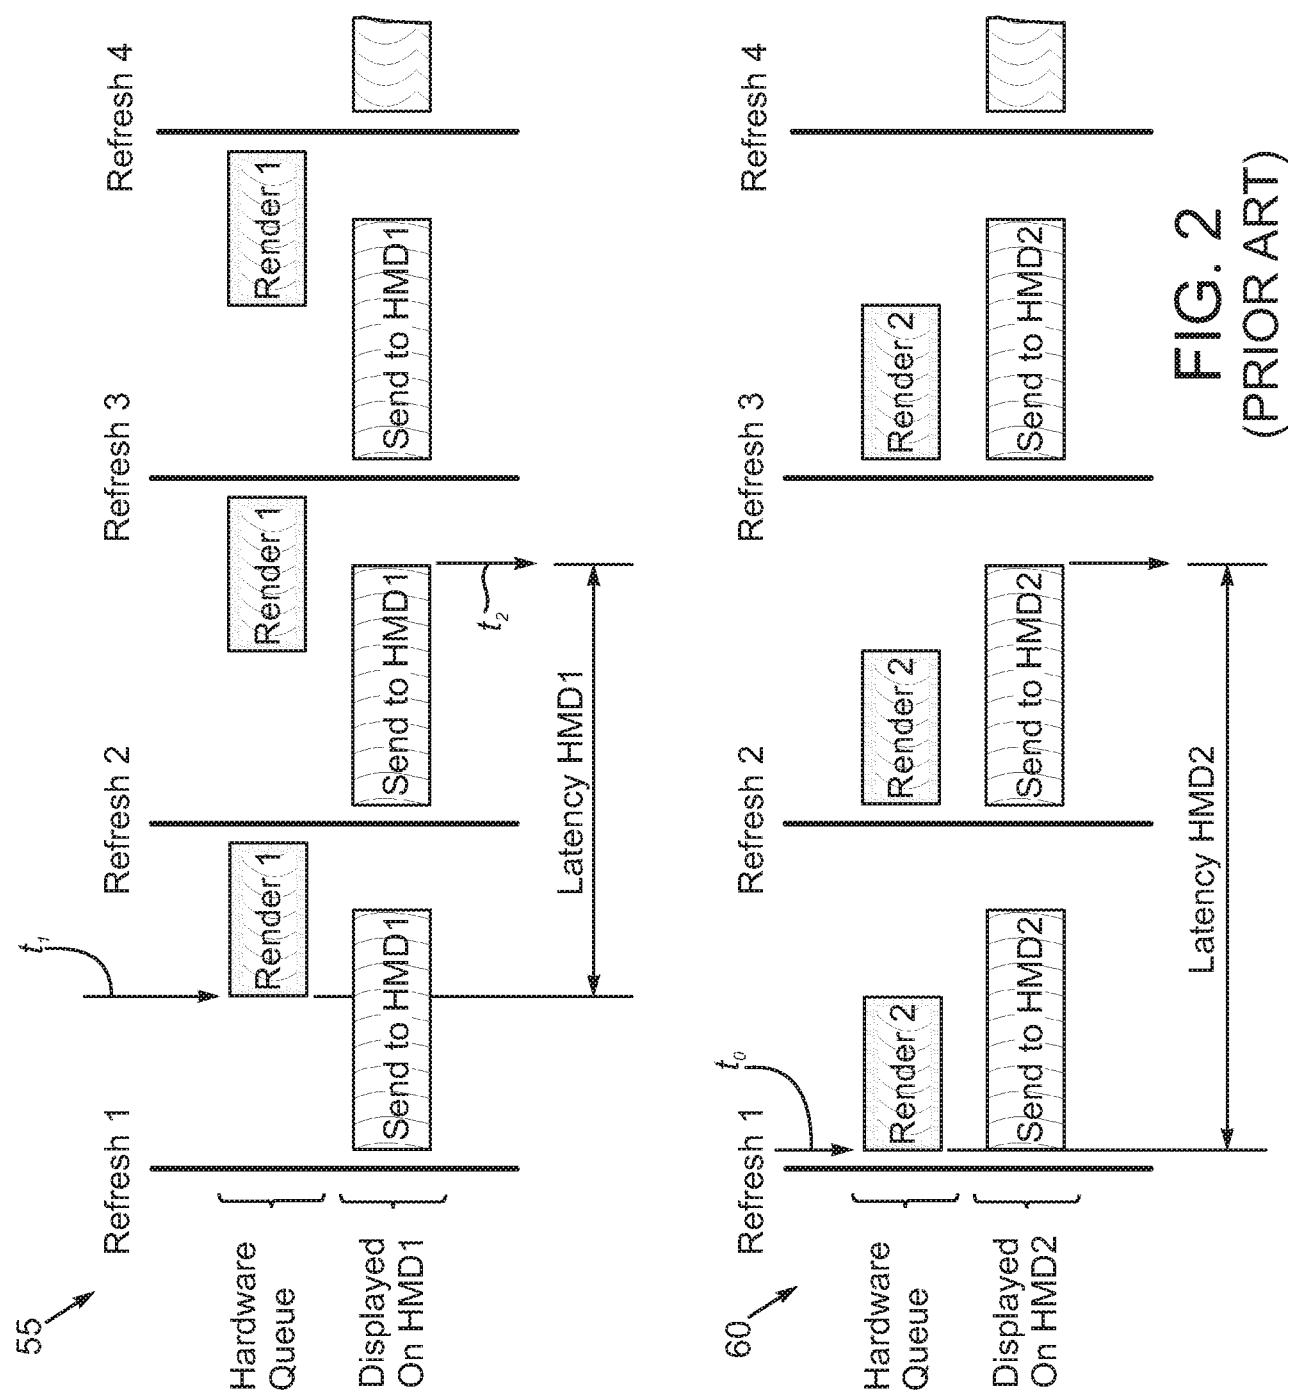


FIG. 2
(PRIOR ART)

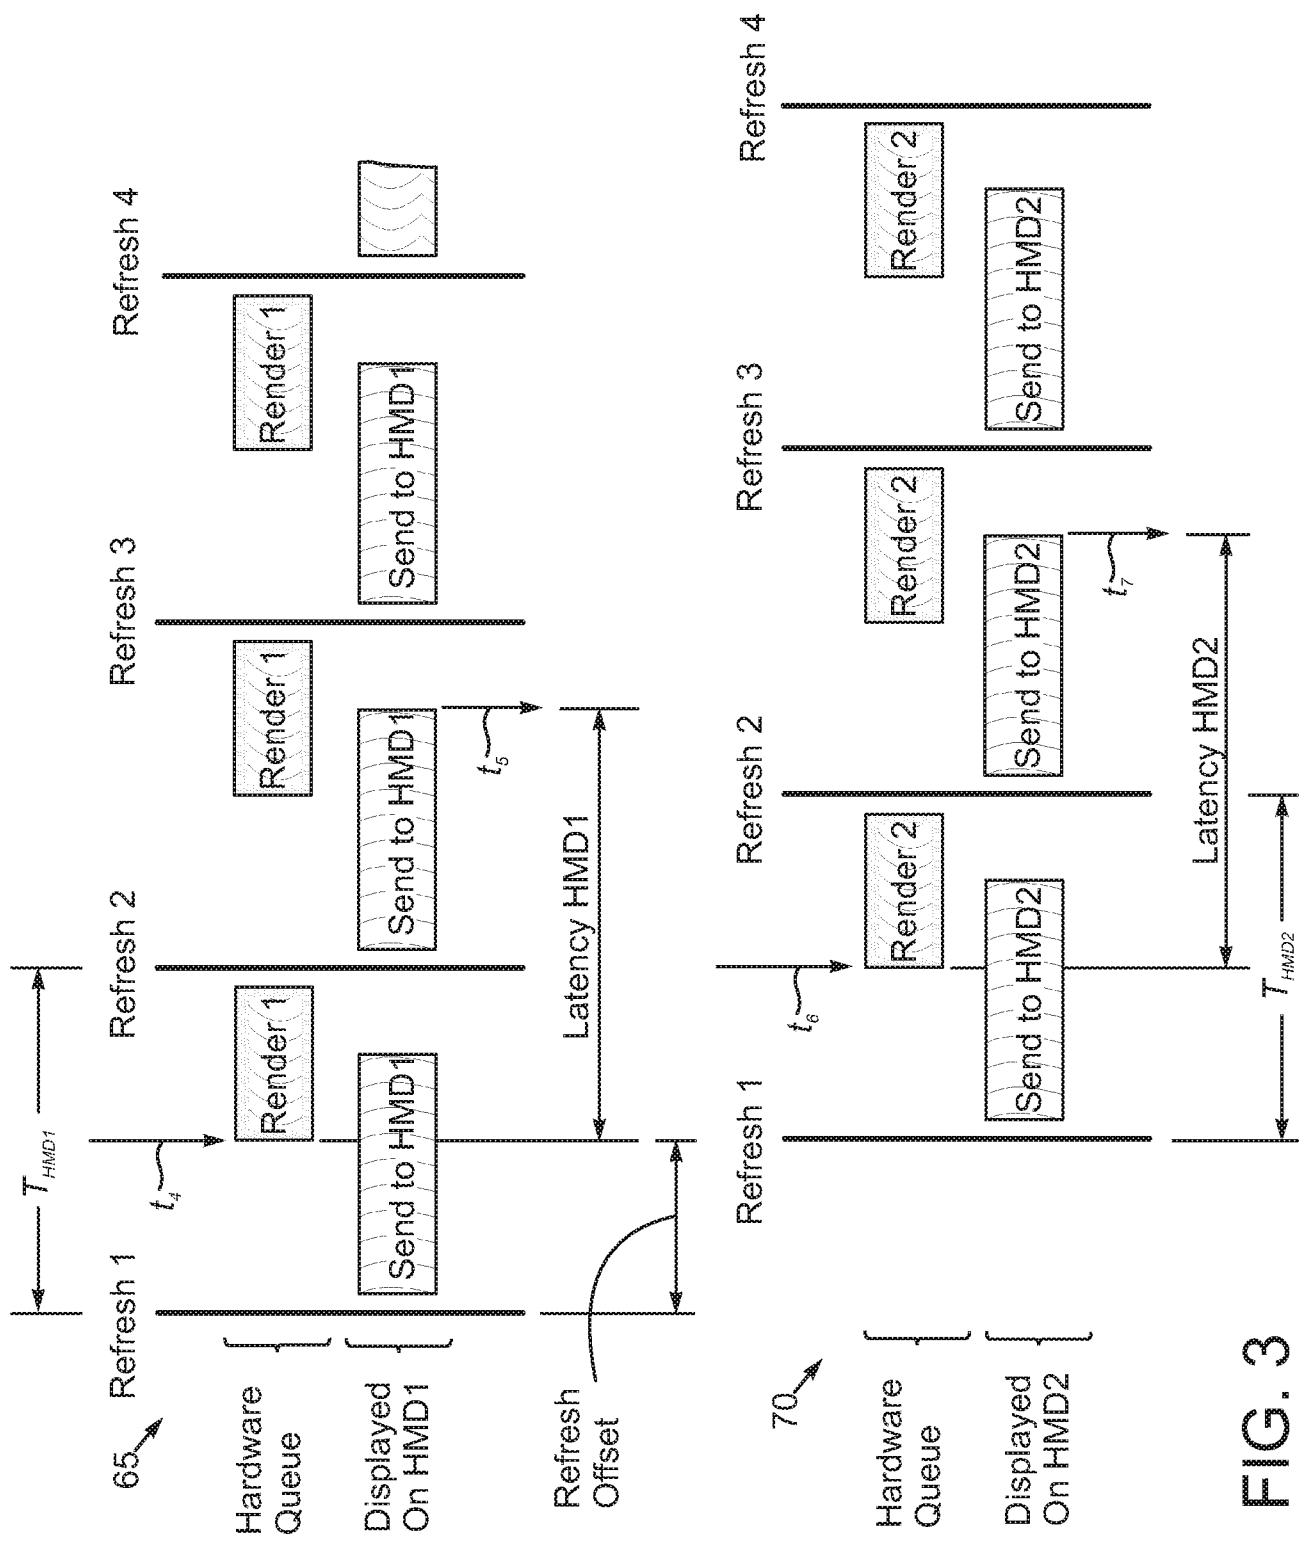


FIG. 3

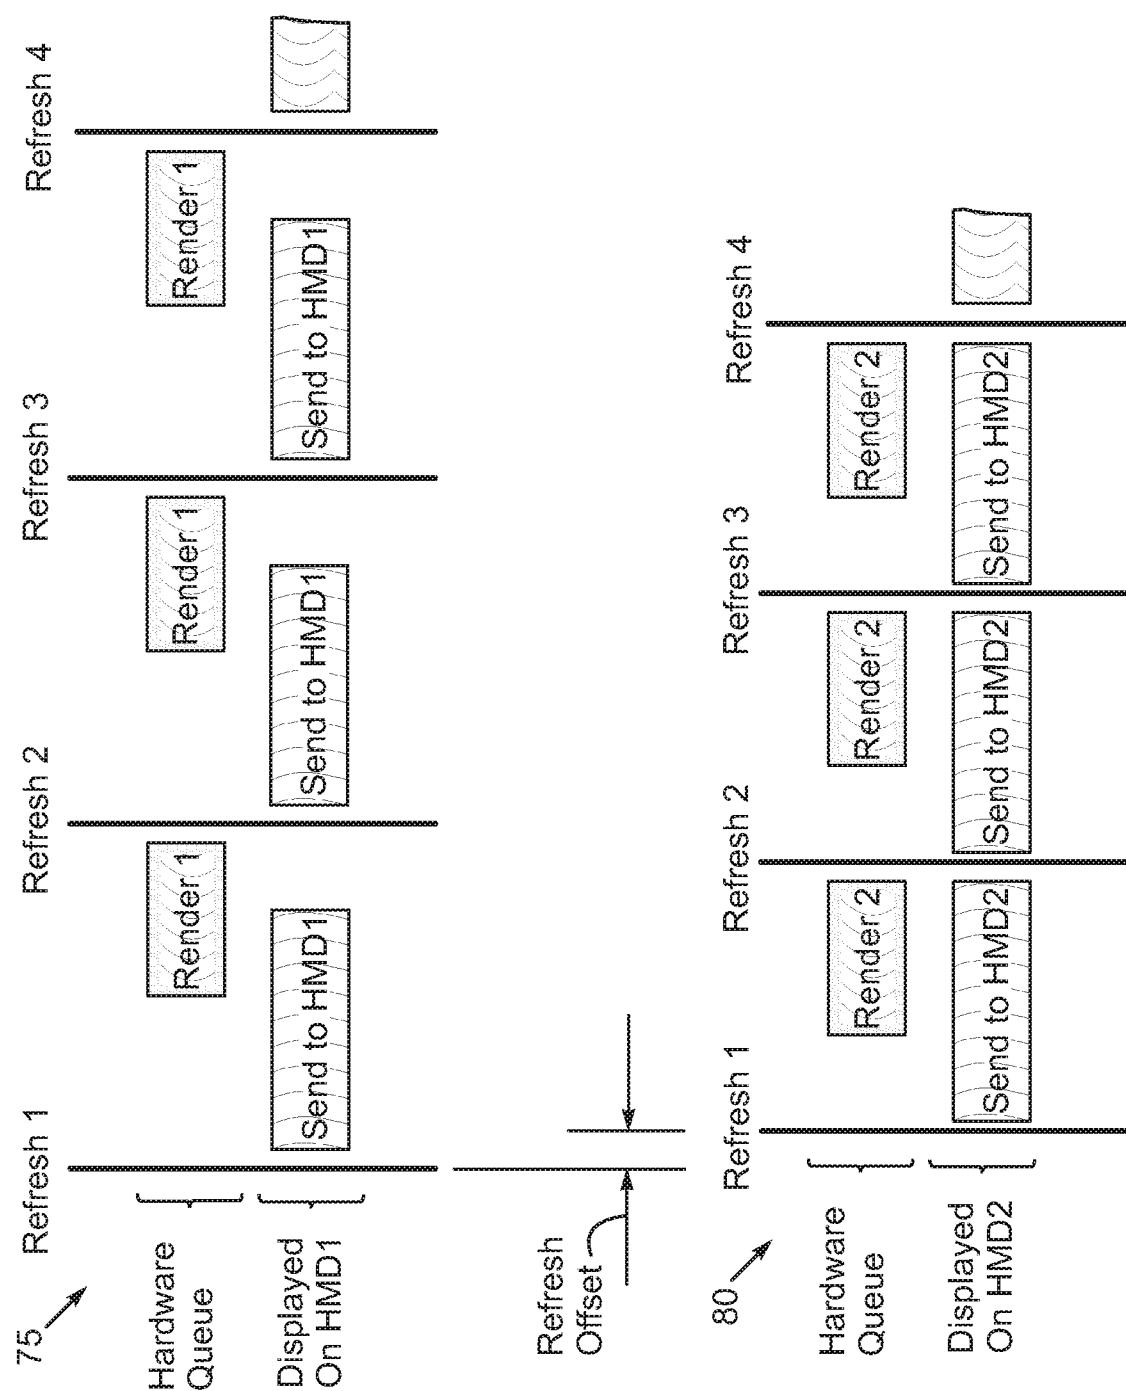


FIG. 4

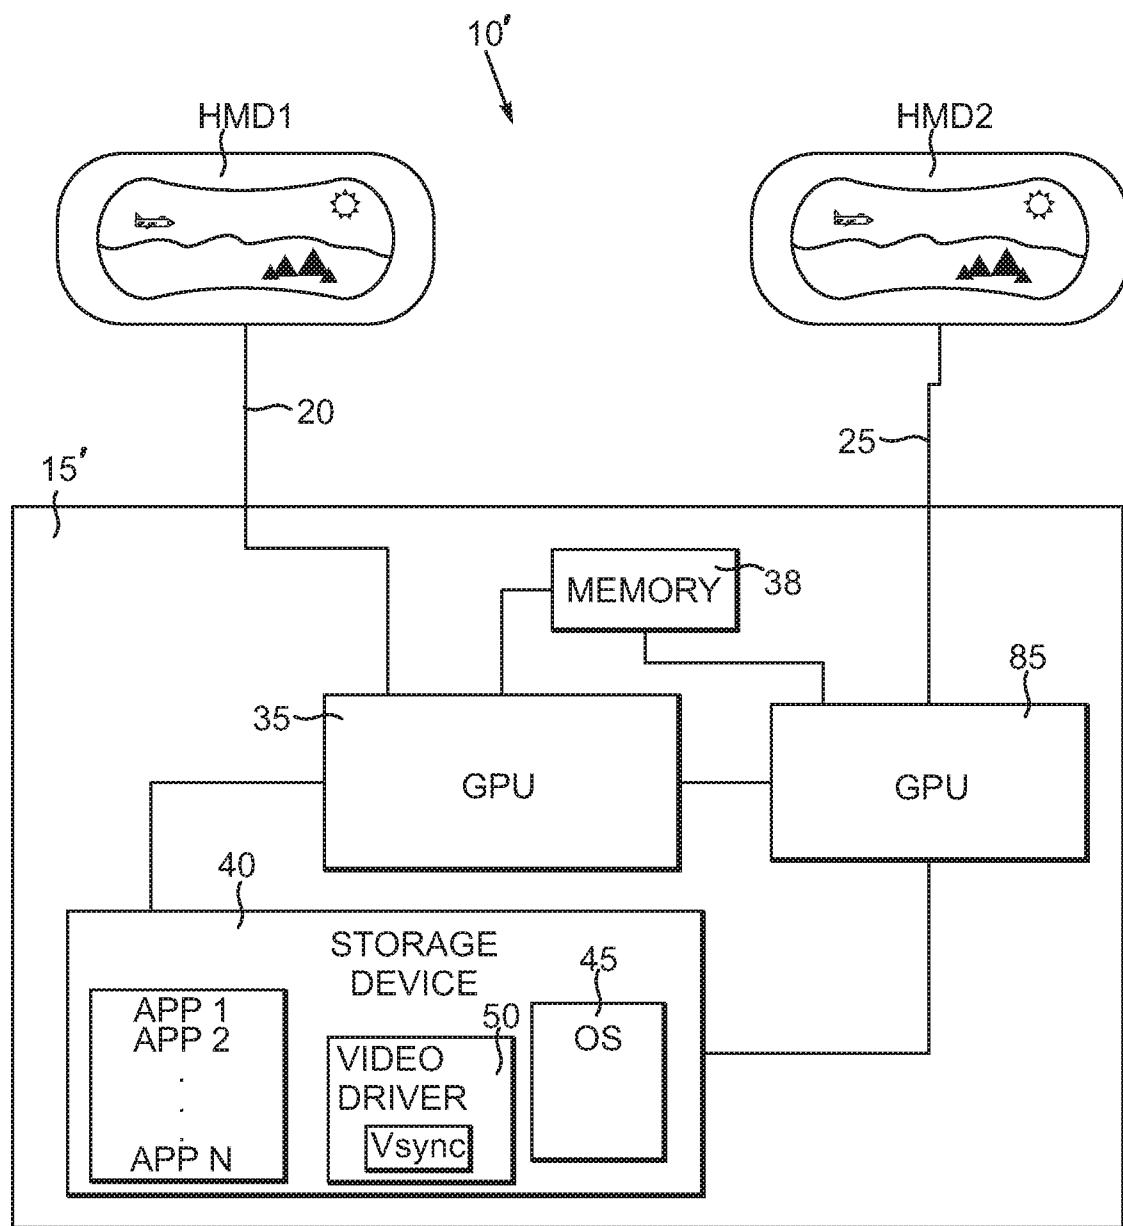


FIG. 5

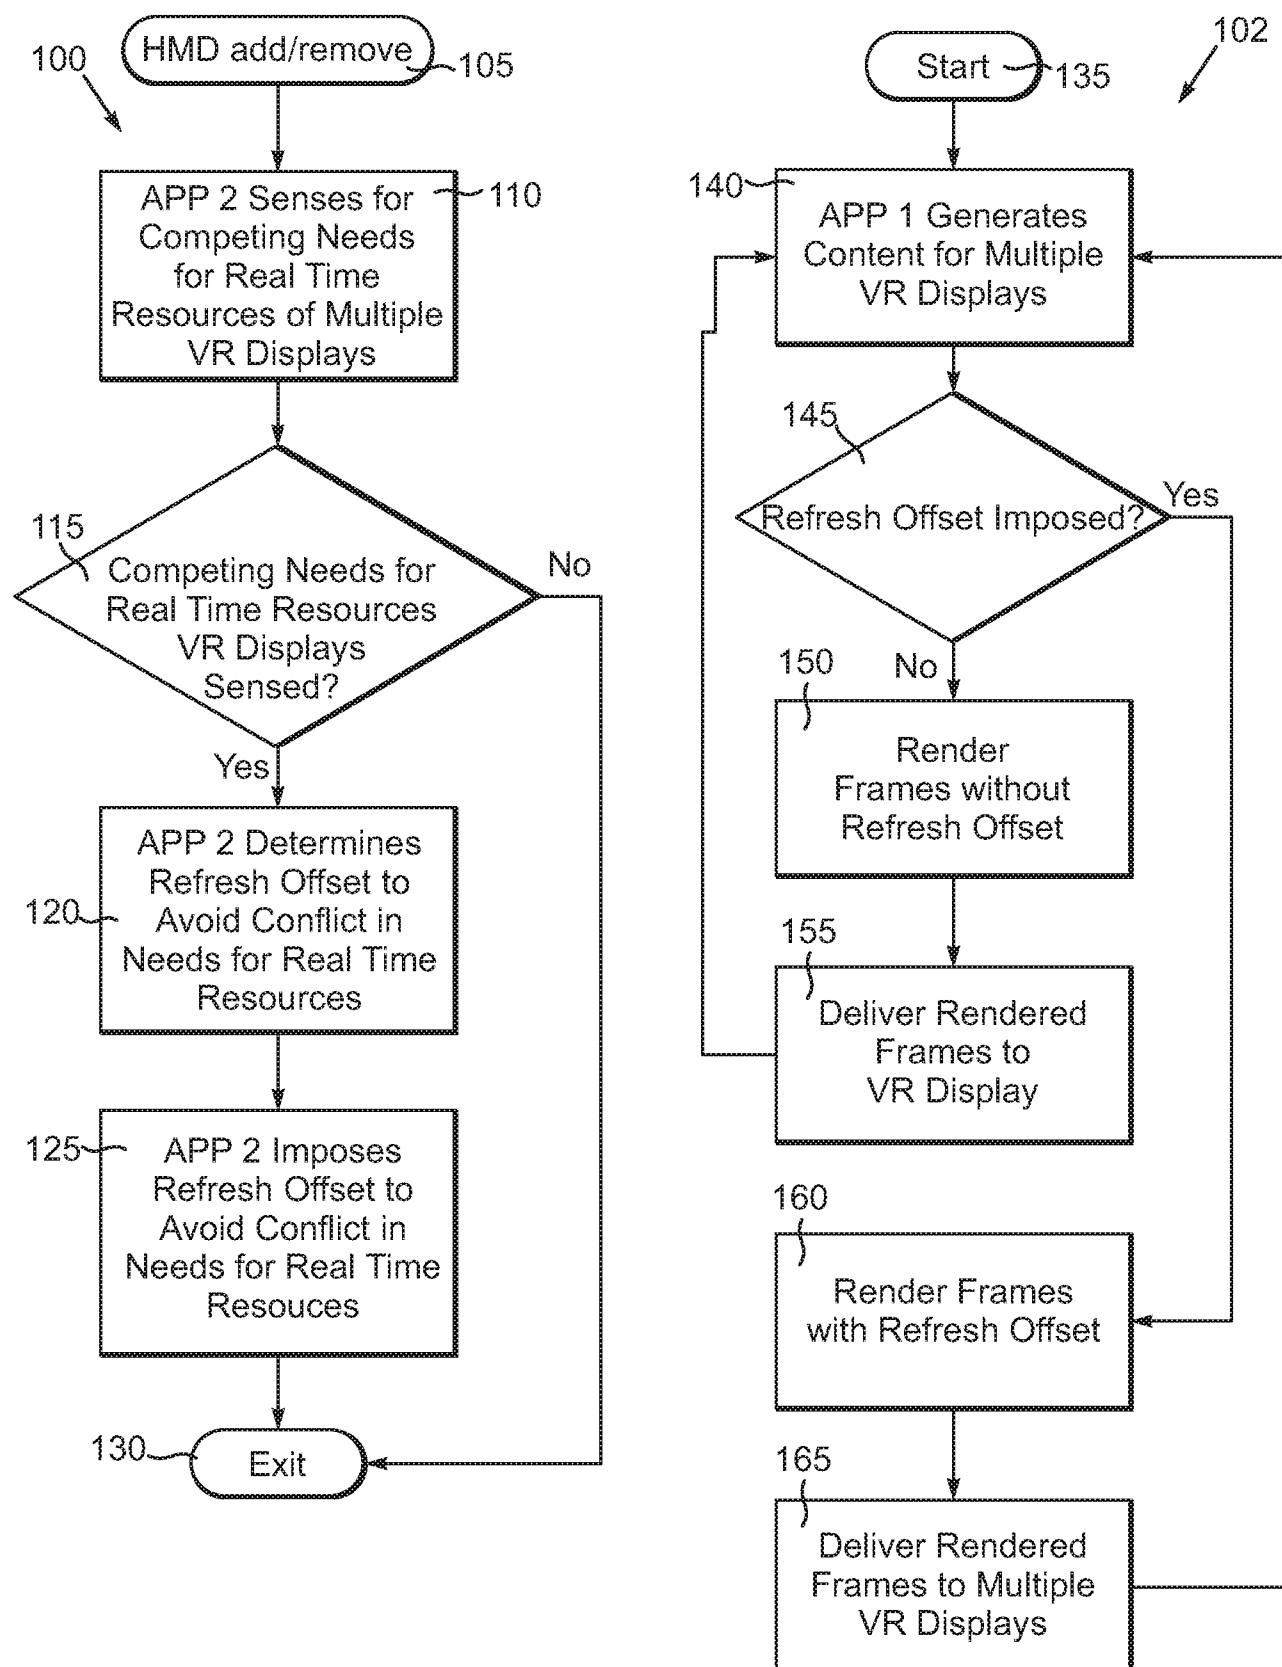


FIG. 6

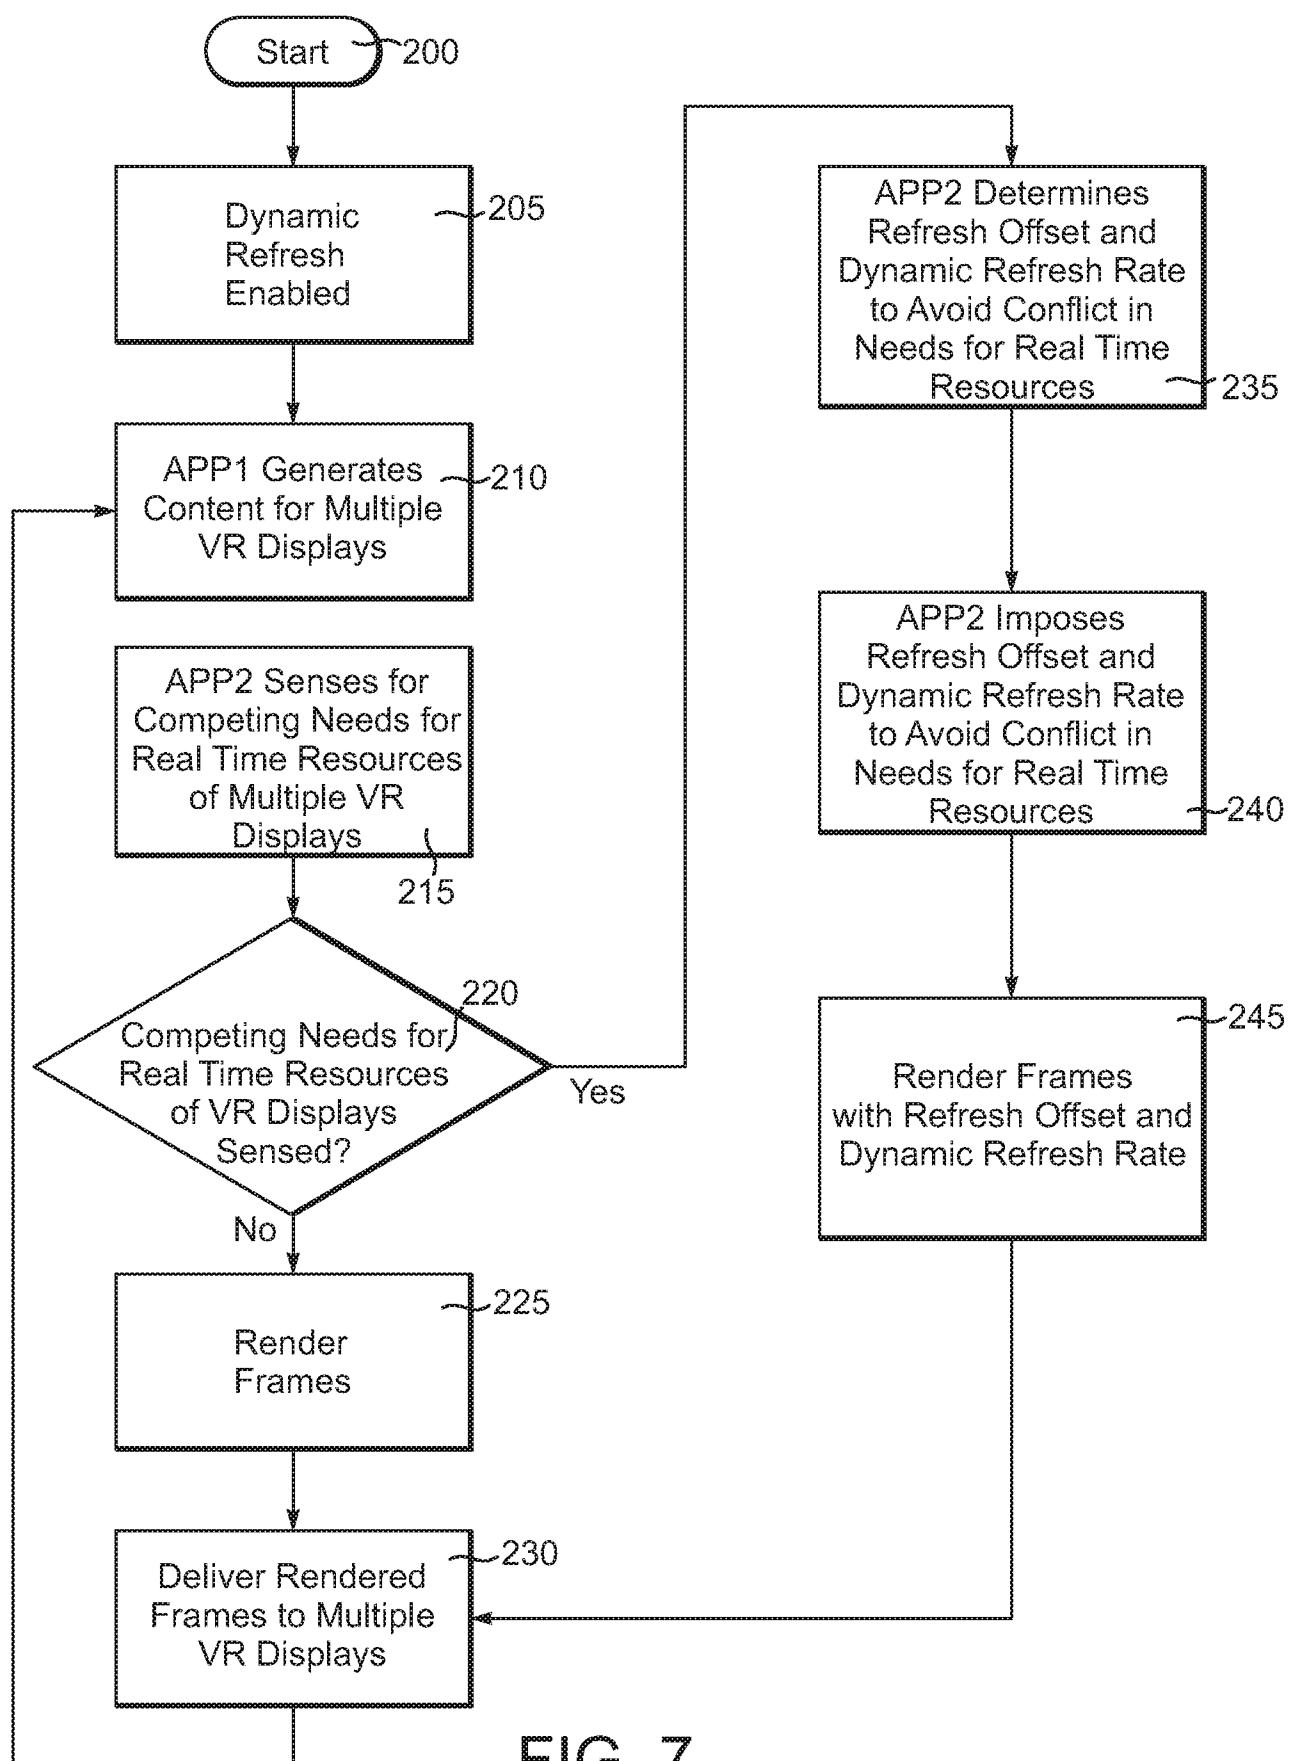


FIG. 7

INTERNATIONAL SEARCH REPORT

International application No.
PCT/IB2018/052735

A. CLASSIFICATION OF SUBJECT MATTER

IPC: **H04N 13/144** (2018.01), **A63F 13/52** (2014.01), **G02B 27/01** (2006.01), **G09G 5/14** (2006.01)

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)

IPC: **H04N 13/144** (2018.01), **A63F 13/52** (2014.01), **G02B 27/01** (2006.01), **G09G 5/14** (2006.01)

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

Electronic database(s) consulted during the international search (name of database(s) and, where practicable, search terms used)

DATABASES : QUESTEL ORBIT, USPTO WEST, GOOGLE PATENT

KEY WORD SEARCH: Synchronizing multiple head-mounted displays, refresh rate, computing resources, asynchronous time warp, virtual reality displays, headsets, goggles, dynamic refresh displays, GPU, delivering video frame, delivering video game, game console, playing games,

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages				Relevant to claim No.
Y	US2016/0093108 A1 Whole document	Mao et al.	31 March 2016	(31.03.2016)	1-20
Y	US2014/0152676 A1 Whole document	Rohn et al.	5 June 2014	(05.06.2014)	1-20
A	US2017/0011681 A1 Whole document	Bathiche et al.	12 January 2017	(12.01.2017)	1-20
A	US2016/0189429 A1 Whole document	Mallinson	30 June 2016	(30.06.2016)	1-20
A	US2017/0076503 A1 Whole document	Tamaoki et al	16 March 2017	(16.03.2017)	1-20

Further documents are listed in the continuation of Box C.

See patent family annex.

* "A" "E" "L" "O" "P"	Special categories of cited documents: document defining the general state of the art which is not considered to be of particular relevance earlier application or patent but published on or after the international filing date 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) document referring to an oral disclosure, use, exhibition or other means document published prior to the international filing date but later than the priority date claimed	"T" "X" "Y" "&"	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 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 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
12 July 2018 (12-07-2018)

Date of mailing of the international search report
19 July 2018 (19-07-2018)

Name and mailing address of the ISA/CA
Canadian Intellectual Property Office
Place du Portage I, C114 - 1st Floor, Box PCT
50 Victoria Street
Gatineau, Quebec K1A 0C9
Facsimile No.: 819-953-2476

Authorized officer

Adel El Hamad (819) 639-3028

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/IB2018/052735

Patent Document Cited in Search Report	Publication Date	Patent Family Member(s)	Publication Date
US2016093108A1	31 March 2016 (31-03-2016)	US2016093108A1 US9818225B2 CN106716306A EP3201731A1 WO2016053906A1	31 March 2016 (31-03-2016) 14 November 2017 (14-11-2017) 24 May 2017 (24-05-2017) 09 August 2017 (09-08-2017) 07 April 2016 (07-04-2016)
US2014152676A1	05 June 2014 (05-06-2014)	US2014152676A1 CN105027563A EP2926555A1 JP2016509245A KR20150091474A WO2014085788A1	05 June 2014 (05-06-2014) 04 November 2015 (04-11-2015) 07 October 2015 (07-10-2015) 24 March 2016 (24-03-2016) 11 August 2015 (11-08-2015) 05 June 2014 (05-06-2014)
US2017011681A1	12 January 2017 (12-01-2017)	US2017011681A1 US2016180762A1	12 January 2017 (12-01-2017) 23 June 2016 (23-06-2016)
US2016189429A1	30 June 2016 (30-06-2016)	US2016189429A1 US9824498B2 CN107003512A EP3241207A2 JP2018503114A US2018082484A1 WO2016109139A2 WO2016109139A3	30 June 2016 (30-06-2016) 21 November 2017 (21-11-2017) 01 August 2017 (01-08-2017) 08 November 2017 (08-11-2017) 01 February 2018 (01-02-2018) 22 March 2018 (22-03-2018) 07 July 2016 (07-07-2016) 25 August 2016 (25-08-2016)
US2017076503A1	16 March 2017 (16-03-2017)	US2017076503A1 JP2017058971A	16 March 2017 (16-03-2017) 23 March 2017 (23-03-2017)