



- (51) **International Patent Classification:**
H04N 5/76 (2006.01) *H04N 5/00* (2011.01)
- (21) **International Application Number:**
PCT/US2015/031933
- (22) **International Filing Date:**
21 May 2015 (21.05.2015)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (71) **Applicant:** HEWLETT PACKARD ENTERPRISE DEVELOPMENT LP [US/US]; 11445 Compaq Center Drive West, Houston, TX 77070 (US).
- (72) **Inventors:** EMERSON, Theodore F.; 11445 Compaq Center Drive West, Houston, Texas 77070 (US). HEINRICH, David F.; 11445 Compaq Center Drive West, Houston, Texas 77070 (US). CHIN, Kenneth T.; 11445 Compaq Center Drive West, Houston, Texas 77070 (US).
- (74) **Agents:** SHOOKMAN, Jeb et al.; Hewlett Packard Enterprise, 3404 E. Harmony Road, Mail Stop 79, Fort Collins, CO 80528 (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, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, 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))

[Continued on next page]

- (54) **Title:** VIDEO MANAGEMENT FOR COMPUTE NODES




FIG. 2

- (57) **Abstract:** An apparatus includes a plurality of compute nodes and a baseboard management controller that is shared by the plurality of compute nodes to manage video for the compute nodes. The baseboard management controller includes video controllers that are associated with the plurality of compute nodes and at least one resource that is shared by the video controllers.



Published:

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

- 1 -

VIDEO MANAGEMENT FOR COMPUTE NODES

Background

[0001] A typical distributed computer system includes multiple hosts, or compute nodes. For a given computer system, the compute nodes may be formed from physically separate machines, from host partitions within a single physical machine or a combination thereof. An example of a multiple host (or node) distributed computer system is a distributed blade computer system. In this regard, in the blade server architecture, compute nodes are formed from "blades," which are connected to a common power supply. Blades may be mounted in a rack, which includes a power supply and an interconnect structure that is configured to provide remote access to the blades.

- 2 -

Brief Description of the Drawings

[0002] Figs. 1 and 3 are schematic diagrams of multiple host computer systems according to example implementations.

[0003] Fig. 2 is a schematic diagram of a subsystem of a management subsystem for a multiple host computer system according to an example implementation.

[0004] Fig. 4 is a schematic diagram of a management controller of Fig. 3 according to an example implementation.

[0005] Fig. 5 is a flow diagram depicting a technique to communicate refresh data to multiple video controllers that share a video memory according to an example implementation.

Detailed Description

[0006] Embedded management subsystems are ever-increasingly being used in enterprise computing systems. In addition to providing management functions, such as instrumentation and virtual presence, the management subsystem may provide legacy input/output (I/O) for the managed server. Examples of legacy I/O includes serial ports, keyboard/mouse interfaces, flash read only memory (ROM) access, and video. In accordance with example systems and techniques that are disclosed herein, a single management subsystem provides management and legacy I/O services for a plurality of independent compute nodes. Moreover, in accordance with systems and techniques that are disclosed herein, a single management subsystem provides video subsystem I/O for multiple compute nodes.

[0007] As a more specific example, Fig. 1 depicts a computer system 100 in accordance with example implementations. In general, the system 100 is a multiple host system, which includes multiple hosts, called "compute nodes 110" herein. In general, a compute node 110 contains one or more processors (central processing units (CPUs), memory, network interfaces, bridges, and so forth). The compute nodes 110 may be separate physical machines, host partitions formed on a physical machine, or any combination thereof, depending on the particular implementation. As a more specific example, in accordance with some implementations, the compute nodes 110 may be formed in a four socket machine, which is partitioned into two independent two socket systems; and each independent two socket system may form a corresponding compute node 110. It is noted that the system 100 may have more or fewer than two compute nodes 110, in accordance with further example implementations. The compute node 110 includes machine executable instructions, or "software," which when executed by the CPU(s) cause the CPU(s) to provide one or more software components, such as an operating system, application, device drivers, and so forth. In general, the compute nodes 110 are individually-operated computer systems that are coupled to each other through an infrastructure, such as a network fabric (not shown).

[0008] The system 100 of Fig. 1 includes a baseboard management controller 150, which contains video controllers 154 that are associated with the compute nodes

- 4 -

110. As an example, in accordance with some implementations, each video controller 154 may be uniquely associated with a given compute node 110 for purposes of providing a compatible video subsystem for the node 110. The baseboard management controller 150 further includes one or multiple resources 156, which are shared by the video controllers 154. As examples, the resource(s) 156 may include a video memory, a digital-to-analog converter (DAC), a phase locked loop (PLL), and so forth.

[0009] In accordance with example implementations, the video controller 154 may be a video graphics adapter (VGA) or super VGA (SVGA) graphics adapter video subsystem, which allows each compute node 110 to execute an existing operating system without modification for the video services that are provided by the baseboard management controller 150. As a more specific example, the video controller 154 may be an SVGA controller, which is register compatible with the IBM VGA adapter and offers extended modes and capabilities through driver or basic input/output system (BIOS) enhancements. The baseboard management controller 150, in accordance with example implementations, may provide remote keyboard/video/mouse (KVM)-based video redirection for each compute node 110. Moreover, in accordance with example implementations, the resource 156 may include a single dynamic random access memory (DRAM) device that is used by the video controllers 154, which allows the number of compute nodes 110 to scale without imposing additional costly memory requirements on the memory subsystem of the baseboard management controller 150. Moreover, in accordance with example implementations, the number of compute nodes 110 may scale without imposing multiple instances of area intensive I/O macros. As described herein, this allows multiple compute nodes 110 to share a single PLL (a resource 156) and a single set of external display interfaces.

[0010] Referring to Fig. 2, in accordance with an example implementation, a management subsystem 200 for a multiple host computer system includes the video controllers 154, a memory 240 that is shared by the video controllers 154 and which contains memory spaces 244 (or "frame buffers"), which are associated with and used by the corresponding video controllers 154. In this regard, in accordance with

- 5 -

example implementations, a given video controller 154 is uniquely assigned to a memory space 244 and uses the memory space 244 for purposes of rendering its graphic images.

[0011] The video controllers 154, in accordance with example implementations, are "single host" video controllers (single host VGA adapters, for example), which are configured to access the same address space. To accommodate the use of such single host video controllers 154, the subsystem 200 includes an address translator 220. In this manner, the address translator 220 receives first requests 210 from the video controllers 154, where the first requests 210 are directed to the same address space. The address translator 220 converts the first requests 210 into corresponding second requests 224, which allow the video controllers 154 to share the memory 240. In this regard, in accordance with example implementations, the address translator 220 converts the first requests 210 into second requests 224 that are directed to the individual memory spaces 244 and serializes the second requests 224 for purposes of providing the requests 224 to a memory controller 230 (a shared resource). The memory controller 230 may then process the second requests 224 for purposes of communicating graphics rendering traffic with the memory 240.

[0012] Referring to Fig. 3, as a more specific example, in accordance with some implementations, a multiple host computer system 300 includes N compute nodes 310 (compute nodes 310-1 and 310-N, being depicted as examples in Fig. 3). In general, the compute node 310 is a physical machine that is made up of actual hardware 320 and actual machine executable instructions 340, or "software." As an example, the hardware 320 may include one or multiple central processing units (CPUs) 322, a memory 324 and a memory I/O hub, or south bridge 326.

[0013] The memory 324, in general, is a non-transitory storage medium, which may be formed from one or more of the following: semiconductor storage devices, optical storage devices, magnetic storage devices, phase change memory devices, memristors, and so forth. In general, the memory 324 stores program instructions, data, data structures, and so forth for its compute node 310.

[0014] The south bridge 326, in general, may form a communication bridge for the compute node 310 for purposes of communicating (over a communication link 352,

- 6 -

such as a Peripheral Component Interconnect-Express (PCI-E) bus, for example) with a management controller 360 (a baseboard management controller, for example), which is used by all of the compute nodes 110, as described herein. In accordance with example implementations, the south bridge 326 communicates drawing commands and data from the compute node 310 to an associated video controller 364 of the management controller 360.

[0015] In accordance with example implementations, the machine executable instructions 340 may include instructions that are executed by the CPU(s) 322, such as instructions that are associated with an operating system 342, one or multiple applications 346, one or multiple device drivers 348, and so forth. In accordance with example implementations, a given driver 348 may be a video device driver, which constructs drawing commands and transfers data to/from the video controller 364 associated with the compute node 310.

[0016] In accordance with example implementations, the management controller 360 contains two or more video controllers 364. The video controller 364 may be a "single host" video controller, and each of the video controllers 364 may be assigned to different compute node 310. As an example, in accordance with example implementations, the video controller 364 may be a VGA adapter, and the video controllers 364 may be identical. As also depicted in Fig. 3, in general, the management controller 360 may include a volatile memory 366, which may be a shared resource that is used by the video controllers 364. In this manner, in accordance with example implementations, the video controllers 364 may communicate data to/from the volatile memory 366 via a memory bus 350, and volatile memory 366 may contain frame buffers for the associated video controllers 364. Moreover, the volatile memory 366 may store one or more frame buffer contexts as well as code and data for the management controller 360. The management controller 360 may also access a non-volatile memory 368 (via a memory bus 351) for purposes of storing, for example, machine executable instructions, or program instructions, for the management controller 360.

[0017] As also shown in Fig. 3, the management controller 360 may be coupled via a network connection 370 (an Ethernet connection, for example) to a management

- 7 -

network, which may include one or more computers (not shown) that remotely manage the compute nodes 310.

[0018] Referring to Fig. 4 in conjunction with Fig. 3, in accordance with example implementations, the management controller 360 contains two or more video controllers 364. The video controllers 364 are operationally coupled to the compute nodes 310 through corresponding endpoint controllers 440, such as Peripheral Component Interconnect-Express (PCI-E) endpoint controllers, for example. The endpoint controller 440 allows the operating system 342 running on the corresponding compute node 310 to discover, configure, and control its own set of video controller hardware, in accordance with example implementations. As an example, upon discovery of its associated video controller 364, the operating system 342 may load a corresponding display driver to facilitate the rendering of images to either a local display (not shown) or a remote display (not shown) using KVM redirection hardware 490.

[0019] In accordance with example implementations, the video controller 364 contains two memory interfaces: a first memory interface 402 that provides bi-directional (read/write) access to the volatile memory 366 (and a corresponding frame buffer 430) for purposes of rendering display commands that are issued by the operating system driver 348; and a second memory interface 403, which provides read-only access to the rendered frame buffer 430 so that image may be represented on an analog or digital display interface. The second memory interface 403 may be a relatively intensive bandwidth interface, which may read the entire frame buffer 430 (several megabytes in size, for example) at a relatively fast rate (60-85 times per second, for example).

[0020] More specifically, in accordance with example implementations, the first memory interface 402 for each video controller 364 is coupled to a cycle funnel 410. In general, the cycle funnel 410 is an address translator, which is constructed to allow unified serialized access to the shared memory controller 420 resource. The cycle funnel 410 receives each request from the memory interface 402, which is a request directed to rendering traffic 404. The rendering traffic 404, in turn, is directed to the same memory space for each video controller 364, in accordance

- 8 -

with example implementations. The cycle funnel 410 takes each request received from the video controllers 364, translates the address into a unified memory model and re-issues the requests in a serialized fashion to the memory controller 420, as represented by a unified rendering traffic 412 in Fig. 4.

[0021] The cycle funnel 410 is also responsible for returning read information to the appropriate requester. The cycle funnel 410 may, in accordance with example implementation, insert wait states on each request to ensure that each video controller 364 receives an equal amount of rendering traffic bandwidth.

[0022] In accordance with example implementations, the second memory interface 403 of the video controller 364 receives display refresh traffic 450. The second memory interfaces 403 are operationally coupled to a display refresh multiplexer 460. Unlike the cycle funnel 410, the display refresh multiplexer 460 allows a selected subset of one or more video controllers 364 to be coupled to the memory 366, and this selected subset of one or more video controllers 364 furnish video output via corresponding digital video output (DVO) signals 481. As depicted in Fig. 4, the display refresh multiplexer 460 communicates selected refresh traffic data 470 to/from the memory controller 420. The video controller(s) 364 may be selected using a compute node selection signal that is received via terminal(s) 461 of the display refresh multiplexer 460. The selection of the subset of video controller(s) 364 that are actually being used to provide video output(s) allows the display refresh multiplexer 460 to satisfy real-time requests without overwhelming or oversaturating the shared memory resource, in accordance with example implementations. Accordingly, the selected subset of video controller(s) 364 that are allowed to pass request information through the display refresh multiplexer 460 produce corresponding video output stream(s) that are routed to the appropriate display device(s), whereas requests for other, non-selected video controller(s) 364 are not processed.

[0023] In accordance with example implementations, the display refresh multiplexer 460 provides placebo, or placeholder, data to the non-selected video controller(s) 364 prevent the display engine(s) of the video controller(s) 364 from experiencing side effect(s) due to the engine(s) being unconnected. Using a multiplexer 482 that

- 9 -

multiplexes digital video outputs from the video controllers 364, the video streams from the video controller(s) 364 receiving placebo display refresh information are ignored, or not selected by the multiplexer 482 (as controlled by control terminal(s) 483 of the multiplexer 482, which receives the compute node selection signal).

[0024] For the example implementation of Fig. 4, a single video controller 364 is selected to provide a digital video output signal s(called "DVO" in Fig. 4), and the multiplexer 482 selects the single video controller 364 for purposes of providing the output stream from the selected video controller 364 as a video output for the management controller 360.

[0025] Thus, the display refresh multiplexer 460 ensures that each video controller 364 receives valid interface response signals regardless of its connected state (i.e., regardless of whether a given video controller 364 is providing a real DVO output stream or not). The display refresh multiplexer 460 further ensures that any outstanding requests from a connected video controller 364 are satisfied before allowing a configuration switch of the connected video controller(s) 364.

[0026] The cycle funnel 410 allows each compute node 310 to separately and simultaneously render images into the frame buffers 430 that, in accordance with example implementations, are uniquely assigned to the compute nodes 310 (and associated video controllers 364). The multiplexer 482 allows an administrator to selectively display output from one or multiple sources. In this regard, in accordance with some implementations, management control firmware may perform periodic selection of the "connected" video controller so that video redirection logic may take a "snapshot" of the video output from multiple node instances at a reduced refresh rate. A link manager 492 of the management controller 360 establishes and trains the source and sync, sets the speed of the link, and so forth. Firmware that is executed on the management controller 360 may, in accordance with example implementations, periodically control terminals 483, 499 and 461 for purposes of obtaining snapshots of the video output from multiple node instances. Firmware executed on the management controller may also perform the job of unifying the image or creating a web page, in accordance with example implementations.

- 10 -

[0027] As depicted in Fig. 4, the management controller 360 may include a digital interface 494 (containing interfaces, such as DisplayPort, HDMI, DVI, and so forth) to provide a digital video output 495 (and a DAC 496 to provide an analog output 497. The management controller 360 may also include KVM redirection logic 490.

[0028] In accordance with example implementations, for purposes of providing a relatively stable output image, the video controller 364 receives a relatively precise and relatively low jitter clock signal. Moreover, the video controller 364, regardless of its "connected" status, uses a clock input of varying frequency, based on its display mode. In accordance with example implementations, the management controller includes a separate phase locked loop (PLL) for each video controller 364.

[0029] In accordance with further example implementations, for purposes of conserving die area and power, the management controller 360 includes a single PLL 447 and a pseudo-PLL 443 for each video controller 364, as depicted in Fig. 4. In accordance with example implementations, the pseudo PLL 443 is a non-PLL-based frequency divider (a clocked register, for example). In this regard, as shown in Fig. 4, in accordance with example implementations, the PLL 447 receives input signals or control lines 451 for purposes of selecting a clock signal that is provided by the PLL 447 to the selected video controller 364 for purposes of the controller 364 providing its output and receiving the display refresh traffic from the display refresh multiplexer 460.

[0030] The other, non-selected video controller(s) 364 receive their clock signal(s) from associated pseudo PLL(s) 443. The pseudo PLL 443 has select inputs 441, which specify the desired frequency that is requested by the video controller 364 for the configured display mode. In this regard, a clock selection circuit 445 of the management controller 360 is controlled by a node select signal (received at control terminal(s) 499) and provides the clock signals to the video controllers 364 (via clock lines 446). The clock selection circuit 445 receives a clock signal from the PLL 447 on an input terminal 453. The clock selection circuit 445 switches between the low jitter, area intensive PLL clock output and a relatively higher jitter but more area efficient pseudo PLL clock output based on the node select signal. This arrangement, in accordance with example implementations, has the advantage of

- 11 -

ensuring that all video controllers 364 are provided a clock signal having the correct frequency, where the valuable, low jitter clocking resources are assigned to video controller(s) 364 that are actively providing output video signals.

[0031] In accordance with example implementations, the volatile memory 366 may also be used for management functions. Thus, as shown in Fig. 4, one or multiple other memory requestors 480 inside the management controller 360 may access the memory 366 by interacting with the memory controller 420. The memory controller 420 ensures exclusivity of the memory resources based on the requestor, in accordance with example implementations.

[0032] Referring to Fig. 5, thus, in accordance with example implementations, a technique 500 includes coupling (block 504) compute nodes to video controllers and selecting (block 508) a first set of one or more of the video controllers. Pursuant to the technique 500, display refresh data is communicated (block 512) from a frame memory to the selected first set of video controller(s) and placeholder data is communicated (block 516) in lieu of refresh data to the remaining video controller(s). Pursuant to the technique 500, video output data from the first set of video controller(s) may then be communicated (block 520) to the video output device.

[0033] While the present techniques have been described with respect to a number of embodiments, it will be appreciated that numerous modifications and variations may be applicable therefrom. It is intended that the appended claims cover all such modifications and variations as fall within the scope of the present techniques.

- 12 -

What is claimed is:

1 1. An apparatus comprising:
2 a plurality of compute nodes;
3 a baseboard management controller shared by the plurality of compute nodes
4 to manage video for the compute nodes, the baseboard management controller
5 comprising video controllers associated with the plurality of compute nodes; and
6 at least one resource shared by the video controllers.

1 2. The apparatus of claim 1, wherein the at least one resource comprises
2 a memory controller and a memory, the baseboard management controller further
3 comprises at least one requestor to further share the memory controller and the
4 memory, and the at least one requestor to provide memory requests associated with
5 instrumentation or virtual presence.

1 3. The apparatus of claim 1, wherein each of the video controllers is
2 associated with each of the compute nodes, and each of the compute nodes
3 discovers, configures and controls its associated video controller.

4 4. The apparatus of claim 1 wherein at least one of the video controllers
5 comprises a single host video graphics adapter (VGA).

- 13 -

1 5. An apparatus comprising:
2 a plurality of video controllers associated with a plurality of compute nodes
3 and to provide first requests associated graphics rendering operations to a first
4 memory address space shared in common by the video controllers;
5 an address translator to receive the first requests and convert the first
6 requests into second requests directed to a plurality of separate memory spaces
7 such that each separate memory space is associated with a different video controller
8 of the plurality of controllers; and
9 a memory controller to receive the second requests and generate memory
10 cycles with a memory in response thereto.

1 6. The apparatus of claim 5, wherein the address translator receives at
2 least some of the first requests in a parallel fashion from at least two video
3 controllers of the plurality of video controllers and provides the at least some first
4 requests serially to the memory controller.

1 7. The apparatus of claim 5, further comprising a refresh multiplexer to
2 receive a set of time overlapping display refresh requests from a group of the video
3 controllers, select one or more of the video controllers, provide data from a rendered
4 frame video memory to the selected one or more video controllers, and provide
5 substitute data to the one or more non-selected video controllers of the group.

1 8. The apparatus of claim 7, wherein each video controller of the group
2 provides video output signals, the apparatus further comprising:
3 a video stream multiplexer to select the video output signal for each selected
4 video controller.

- 14 -

1 9. The apparatus of claim 7, wherein the selected one or more video
2 controllers comprise a first video controller and the non-selected one or more video
3 controllers comprise a second video controller, the apparatus further comprising:
4 a phase locked loop (PLL) to provide a first clock signal;
5 a non-PLL based frequency divider to provide a second clock signal; and
6 a clock selector to, in response to the selection of the first video controller,
7 provide the first clock signal to the first video controller to clock the first video
8 controller and, in response to the non-selection of the second video controller,
9 provide the second clock signal to the second video controller to clock the second
10 video controller.

1 10. The apparatus of claim 5, further comprising at least one memory
2 requestor associated with a compute node management controller to communicate
3 memory requests to the memory controller.

1 11. The apparatus of claim 5, wherein the compute nodes comprise host
2 partitions formed on a physical machine.

1 12. A method comprising:
2 coupling a plurality of compute nodes to a plurality of video controllers;
3 selecting a first set of at least one video controller of the plurality of video
4 controllers, a remaining one or more video controllers of the plurality of video
5 controllers being part of a second set of at least one video controller;
6 communicating display refresh data from a frame memory to the first set of at
7 least one video controller;
8 communicating placeholder data in lieu of display refresh data to the second
9 set of at least one video controller; and
10 communicating video output data from the first set of at least one video
11 controller to a video output device for the plurality of compute nodes.

- 15 -

1 13. The method of claim 12, further comprising:
2 receiving a first set of refresh requests from the first set of at least one video
3 controller;
4 receiving a second set of refresh requests from the second set of at least one
5 video controller;
6 forwarding the first set of refresh requests to a memory controller to retrieve
7 the display refresh data from the frame memory; and
8 preventing the second set of refresh requests from being forwarded to the
9 memory controller.

1 14. The method of claim 12, wherein the first set of at least one video
2 controller comprises multiple video controllers, the method further comprising
3 processing the video output data to combine snapshots from the multiple video
4 controllers into data representing a screen image.

1 15. The method of claim 12, further comprising:
2 communicating graphics rendering data between a frame memory and the
3 first set of at least one video controller; and
4 communicating graphics rendering data between the frame memory and the
5 second set of at least one video controller.




FIG. 1

2 / 5




FIG. 2




FIG. 3




FIG. 4

5 / 5




FIG. 5

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2015/031933**A. CLASSIFICATION OF SUBJECT MATTER****H04N 5/76(2006.01)i, H04N 5/00(2011.01)i**

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)

H04N 5/76; G06F 15/173; G06F 3/14; G06F 13/14; G06T 1/60; G09G 5/377; G06F 15/167; G06T 1/00; G06F 13/12; H04N 5/00

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: video management, a plurality of video controller, address translator, request, display refresh data, placeholder data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2012-0098841 A1 (THEODORE F. EMERSON) 26 April 2012 See paragraphs [0013], [0016]-[0018]; claim 12; and figure 1.	1-4
A		5-15
A	US 2012-0327094 A1 (SZE HAU LOH et al.) 27 December 2012 See paragraphs [0025], [0042]; and figure 4.	1-15
A	US 7966402 B2 (THEODORE F. EMERSON et al.) 21 June 2011 See column 3, lines 37-39; claim 20; and figure 2.	1-15
A	US 2014-0176587 A1 (STMICROELECTRONICS, INC.) 26 June 2014 See paragraph [0018]; claim 5; and figure 7.	1-15
A	US 8922585 B2 (JEFFREY S. AUTOR et al.) 30 December 2014 See column 7, lines 59-64; claim 8; and figure 3.	1-15



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"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

20 April 2016 (20.04.2016)

Date of mailing of the international search report

22 April 2016 (22.04.2016)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsu-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

KIM, KI HO

Telephone No. +82-42-481-8691



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2015/031933

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2012-0098841 A1	26/04/2012	CN 102473079 A CN 102473079 B EP 2454654 A1 EP 2454654 A4 WO 2011-008205 A1	23/05/2012 30/09/2015 23/05/2012 09/01/2013 20/01/2011
US 2012-0327094 A1	27/12/2012	CN 102906718 A EP 2577476 A1 EP 2577476 A4 WO 2011-149480 A1	30/01/2013 10/04/2013 27/08/2014 01/12/2011
US 7966402 B2	21/06/2011	CN 1892586 A CN 1892586 B US 2006-0294254 A1	10/01/2007 01/09/2010 28/12/2006
US 2014-0176587 A1	26/06/2014	EP 0827110 A2 EP 0827110 A3 H1 0145739 A JP 10-145739 A JP 3954164 B2 US 06058459 A US 2002-180743 A1 US 2008-088637 A1 US 2009-201305 A1 US 2010-309211 A1 US 2011-122946 A1 US 2012-007871 A1 US 2013-042072 A1 US 6427194 B1 US 7321368 B2 US 7542045 B2 US 7777753 B2 US 7898548 B2 US 8054315 B2 US 8314808 B2 US 8681164 B2	04/03/1998 26/05/1999 29/05/1998 29/05/1998 08/08/2007 02/05/2000 05/12/2002 17/04/2008 13/08/2009 09/12/2010 26/05/2011 12/01/2012 14/02/2013 30/07/2002 22/01/2008 02/06/2009 17/08/2010 01/03/2011 08/11/2011 20/11/2012 25/03/2014
US 8922585 B2	30/12/2014	US 2008-094413 A1	24/04/2008