



(51) International Patent Classification:

G06T 19/20 (2011.01) H04N 13/04 (2006.01)  
G06T 17/00 (2006.01)

(21) International Application Number:

PCT/US2016/047425

(22) International Filing Date:

17 August 2016 (17.08.2016)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/206,765 18 August 2015 (18.08.2015) US

(71) Applicant: **MAGIC LEAP, INC.** [US/US]; 1855 Griffin Road, Suite B454, Dania Beach, FL 33004 (US).

(72) Inventors: **RODRIGUEZ, Jose Felix**; 1855 Griffin Road, Suite B454, Dania Beach, FL 33004 (US). **PEREZ, Ricardo Martinez**; 1855 Griffin Road, Suite B454, Dania Beach, FL 33004 (US).

(74) Agent: **ALTMAN, Daniel, E.**; Knobbe, Martens, Olson & Bear LLP, 2040 Main Street 14th Floor, Irvine, CA 92614 (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 applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

Published:

- with international search report (Art. 21(3))
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

(54) Title: VIRTUAL AND AUGMENTED REALITY SYSTEMS AND METHODS




FIG. 10

(57) Abstract: A virtual or augmented reality display system that controls a display using control information included with the virtual or augmented reality imagery that is intended to be shown on the display. The control information can be used to specify one of multiple possible display depth planes. The control information can also specify pixel shifts within a given depth plane or between depth planes. The system can also enhance head pose measurements from a sensor by using gain factors which vary based upon the user's head pose position within a physiological range of movement.

WO 2017/031246 A1

## VIRTUAL AND AUGMENTED REALITY SYSTEMS AND METHODS

### CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of priority under 35 USC §119(e) of U.S. Provisional Application No. 62/206,765, filed on August 18, 2015, entitled “VIRTUAL AND AUGMENTED REALITY SYSTEMS AND METHODS,” which is incorporated by reference herein in its entirety.

### BACKGROUND

#### Field

[0002] This disclosure relates to virtual and augmented reality imaging and visualization systems.

#### Description of the Related Art

[0003] Modern computing and display technologies have facilitated the development of virtual reality and augmented reality systems. Virtual reality, or “VR,” systems create a simulated environment for a user to experience. This can be done by presenting computer-generated imagery to the user through a display. This imagery creates a sensory experience which immerses the user in the simulated environment. A virtual reality scenario typically involves presentation of only computer-generated imagery rather than also including actual real-world imagery.

[0004] Augmented reality systems generally supplement a real-world environment with simulated elements. For example, augmented reality, or “AR,” systems may provide a user with a view of the surrounding real-world environment via a display. However, computer-generated imagery can also be presented on the display to enhance the real-world environment. This computer-generated imagery can include elements which are contextually-related to the real-world environment. Such elements can include simulated text, images, objects, etc. The simulated elements can often times be interactive in real time. Figure 1 depicts an example augmented reality scene 1 where a user of an AR technology sees a real-world park-like setting 6 featuring people, trees, buildings in the background, and a concrete platform 1120. In addition to these items, computer-generated imagery is also

presented to the user. The computer-generated imagery can include, for example, a robot statue 1110 standing upon the real-world platform 1120, and a cartoon-like avatar character 2 flying by which seems to be a personification of a bumble bee, even though these elements 2, 1110 are not actually present in the real-world environment.

**[0005]** Because the human visual perception system is complex, it is challenging to produce a VR or AR technology that facilitates a comfortable, natural-feeling, rich presentation of virtual image elements amongst other virtual or real-world imagery elements.

#### SUMMARY

**[0006]** In some embodiments, a virtual or augmented reality display system comprises: a display configured to display imagery for a plurality of depth planes; a display controller configured to receive rendered virtual or augmented reality imagery data from a graphics processor, and to control the display based at least in part on control information embedded in the rendered imagery, wherein the embedded control information indicates a shift to apply to at least a portion of the rendered imagery when displaying the imagery.

**[0007]** In some embodiments, the shift alters the displayed position of one or more virtual or augmented reality objects as compared to the position of the one or more objects in the rendered imagery.

**[0008]** In some embodiments, the shift comprises a lateral shift of at least a portion of the imagery by a specified number of pixels within the same depth plane.

**[0009]** In some embodiments, the shift comprises a longitudinal shift of at least a portion of the imagery from one depth plane to another.

**[0010]** In some embodiments, the display controller is further configured to scale at least a portion of the imagery in conjunction with a longitudinal shift from one depth plane to another.

**[0011]** In some embodiments, the shift comprises a longitudinal shift of at least a portion of the imagery from one depth plane to a virtual depth plane, the virtual depth plane comprising a weighted combination of at least two depth planes.

**[0012]** In some embodiments, the shift is based on information regarding a head pose of a user.

[0013] In some embodiments, the shift is performed by the display controller without re-rendering the imagery.

[0014] In some embodiments, a method in a virtual or augmented reality display system comprises: receiving rendered virtual or augmented reality imagery data from a graphics processor; and displaying the imagery for a plurality of depth planes based at least in part on control information embedded in the rendered imagery, wherein the embedded control information indicates a shift to apply to at least a portion of the rendered imagery when displaying the imagery.

[0015] In some embodiments, the method further comprises shifting the displayed position of one or more virtual or augmented reality objects as compared to the position of the one or more objects in the rendered imagery.

[0016] In some embodiments, the method further comprises laterally shifting at least a portion of the imagery by a specified number of pixels within the same depth plane based on the control information.

[0017] In some embodiments, the method further comprises longitudinally shifting at least a portion of the imagery from one depth plane to another based on the control information.

[0018] In some embodiments, the method further comprises scaling at least a portion of the imagery in conjunction with longitudinally shifting the imagery from one depth plane to another.

[0019] In some embodiments, the method further comprises longitudinally shifting at least a portion of the imagery from one depth plane to a virtual depth plane, the virtual depth plane comprising a weighted combination of at least two depth planes.

[0020] In some embodiments, the shift is based on information regarding a head pose of a user.

[0021] In some embodiments, the method further comprises shifting the imagery without re-rendering the imagery.

[0022] In some embodiments, a virtual or augmented reality display system comprises: a display configured to display virtual or augmented reality imagery for a plurality of depth planes, the imagery comprising a series of images made up of rows and columns of

pixel data; a display controller configured to receive the imagery from a graphics processor and to control the display based at least in part on control information embedded in the imagery, wherein the embedded control information comprises depth plane indicator data which indicates at which of the plurality of depth planes to display at least a portion of the imagery.

[0023] In some embodiments, the control information does not alter the number of rows and columns of pixel data in the series of images.

[0024] In some embodiments, the control information comprises a row or column of information substituted for a row or column of pixel data in one or more of the series of images.

[0025] In some embodiments, the control information comprises a row or column of information appended to the pixel data for one or more of the series of images.

[0026] In some embodiments, the pixel data comprises a plurality of color values, and wherein the depth plane indicator data is substituted for one or more bits of at least one of the color values.

[0027] In some embodiments, the depth plane indicator data is substituted for one or more least significant bits of at least one of the color values.

[0028] In some embodiments, the depth plane indicator data is substituted for one or more bits of a blue color value.

[0029] In some embodiments, each pixel comprises depth plane indicator data.

[0030] In some embodiments, the display controller is configured to order the series of images based at least in part on the depth plane indicator data.

[0031] In some embodiments, a method in a virtual or augmented reality display system comprises: receiving virtual or augmented reality imagery from a graphics processor, the imagery comprising a series of images made up of rows and columns of pixel data for a plurality of depth planes; displaying the imagery based at least in part on control information embedded in the imagery, wherein the embedded control information comprises depth plane indicator data which indicates at which of the plurality of depth planes to display at least a portion of the imagery.

[0032] In some embodiments, the control information does not alter the number of rows and columns of pixel data in the series of images.

[0033] In some embodiments, the control information comprises a row or column of information substituted for a row or column of pixel data in one or more of the series of images.

[0034] In some embodiments, the control information comprises a row or column of information appended to the pixel data for one or more of the series of images.

[0035] In some embodiments, the pixel data comprises a plurality of color values, and wherein the depth plane indicator data is substituted for one or more bits of at least one of the color values.

[0036] In some embodiments, the depth plane indicator data is substituted for one or more least significant bits of at least one of the color values.

[0037] In some embodiments, the depth plane indicator data is substituted for one or more bits of a blue color value.

[0038] In some embodiments, each pixel comprises depth plane indicator data.

[0039] In some embodiments, the method further comprises ordering the series of images based at least in part on the depth plane indicator data.

[0040] In some embodiments, a virtual or augmented reality display system comprises: a first sensor configured to provide measurements of a user's head pose over time; and a processor configured to estimate the user's head pose based on at least one head pose measurement and based on at least one calculated predicted head pose, wherein the processor is configured to combine the head pose measurement and the predicted head pose using one or more gain factors, and wherein the one or more gain factors vary based upon the user's head pose position within a physiological range of movement.

[0041] In some embodiments, the first sensor is configured to be head-mounted.

[0042] In some embodiments, the first sensor comprises an inertial measurement unit.

[0043] In some embodiments, the one or more gain factors emphasize the predicted head pose over the head pose measurement when the user's head pose is in a central portion of the physiological range of movement.

[0044] In some embodiments, the one or more gain factors emphasize the predicted head pose over the head pose measurement when the user's head pose is nearer the middle of the physiological range of movement than a limit of the user's physiological range of movement.

[0045] In some embodiments, the one or more gain factors emphasize the head pose measurement over the predicted head pose when the user's head pose approaches a limit of the physiological range of movement.

[0046] In some embodiments, the one or more gain factors emphasize the head pose measurement over the predicted head pose when the user's head pose is nearer a limit of the physiological range of movement than the middle of the physiological range of movement.

[0047] In some embodiments, the first sensor is configured to be head-mounted and further comprising a second sensor configured to be body-mounted, wherein the at least one head pose measurement is determined based on measurements from both the first sensor and the second sensor.

[0048] In some embodiments, the head pose measurement is determined based on a difference between measurements from the first sensor and the second sensor.

[0049] In some embodiments, a method of estimating head pose in a virtual or augmented reality display system comprises: receiving measurements of a user's head pose over time from a first sensor; and estimating, using a processor, the user's head pose based on at least one head pose measurement and based on at least one calculated predicted head pose, wherein estimating the user's head pose comprises combining the head pose measurement and the predicted head pose using one or more gain factors, and wherein the one or more gain factors vary based upon the user's head pose position within a physiological range of movement.

[0050] In some embodiments, the first sensor is configured to be head-mounted and the method further comprises: receiving body orientation measurements from a second sensor configured to be body-mounted; and estimating the user's head pose based on the at least one head pose measurement and based on the at least one calculated predicted head

pose, wherein the at least one head pose measurement is determined based on measurements from both the first sensor and the second sensor.

**[0051]** In some embodiments, a virtual or augmented reality display system comprises: a sensor configured to determine one or more characteristics of the ambient lighting; a processor configured to adjust one or more characteristics of a virtual object based on the one or more characteristics of the ambient lighting; and a display configured to display the virtual object to a user.

**[0052]** In some embodiments, the one or more characteristics of the ambient lighting comprise the brightness of the ambient lighting.

**[0053]** In some embodiments, the one or more characteristics of the ambient lighting comprise the hue of the ambient lighting.

**[0054]** In some embodiments, the one or more characteristics of the virtual object comprise the brightness of the virtual object.

**[0055]** In some embodiments, the one or more characteristics of the virtual object comprise the hue of the virtual object.

**[0056]** In some embodiments, a method in a virtual or augmented reality display system comprises: receiving one or more characteristics of the ambient lighting from a sensor; adjusting, using a processor, one or more characteristics of a virtual object based on the one or more characteristics of the ambient lighting; and displaying the virtual object to a user.

**[0057]** In some embodiments, a virtual or augmented reality display system comprises: a processor configured to compress virtual or augmented reality imagery data, the imagery comprising imagery for multiple depth planes, the processor being configured to compress the imagery data by reducing redundant information between the depth planes of the imagery; a display configured to display the imagery for the plurality of depth planes.

**[0058]** In some embodiments, the imagery for a depth plane is represented in terms of differences with respect to an adjacent depth plane.

**[0059]** In some embodiments, the processor encodes motion of an object between depth planes.



**[0060]** In some embodiments, a method in a virtual or augmented reality display system comprises: compressing virtual or augmented reality imagery data with a processor, the imagery comprising imagery for multiple depth planes, the processor being configured to compress the imagery data by reducing redundant information between the depth planes of the imagery; displaying the imagery for the plurality of depth planes.

**[0061]** In some embodiments, the imagery for a depth plane is represented in terms of differences with respect to an adjacent depth plane.

**[0062]** In some embodiments, the method further comprises encoding motion of an object between depth planes.

**[0063]** In some embodiments, a virtual or augmented reality display system comprises: a display configured to display virtual or augmented reality imagery for a plurality of depth planes; a display controller configured to control the display, wherein the display controller dynamically configures a sub-portion of the display to refresh per display cycle.

**[0064]** In some embodiments, the display comprises a scanning display and the display controller dynamically configures the scanning pattern to skip areas of the display where the imagery need not be refreshed.

**[0065]** In some embodiments, the display cycle comprises a frame of video imagery.

**[0066]** In some embodiments, the display controller increases the video frame rate if the sub-portion of the display to be refreshed decreases in size.

**[0067]** In some embodiments, the display controller decreases the video frame rate if the sub-portion of the display to be refreshed increases in size.

**[0068]** In some embodiments, a method in a virtual or augmented reality display system comprises: displaying virtual or augmented reality imagery for a plurality of depth planes with a display; dynamically configuring a sub-portion of the display to refresh per display cycle.

**[0069]** In some embodiments, the display comprises a scanning display and the method further comprises dynamically configuring the scanning pattern to skip areas of the display where the imagery need not be refreshed.

[0070] In some embodiments, the display cycle comprises a frame of video imagery.

[0071] In some embodiments, the method further comprises increasing the video frame rate if the sub-portion of the display to be refreshed decreases in size.

[0072] In some embodiments, the method further comprises decreasing the video frame rate if the sub-portion of the display to be refreshed increases in size.

[0073] In some embodiments, a virtual or augmented reality display system comprises: a transmitter which transmits an electric or magnetic field that varies in space; a tangible object which allows a user to interact with a virtual object or scene, the tangible object comprising a sensor which detects the electric or magnetic field from the transmitter, wherein measurements from the sensor are used to determine the position or orientation of the tangible object with respect to the transmitter.

[0074] In some embodiments, the transmitter is integrated with a head-mounted portion of the virtual or augmented reality display system.

[0075] In some embodiments, a method in a virtual or augmented reality display system comprises: transmitting an electric or magnetic field that varies in space using a transmitter; detecting the electric or magnetic field using a sensor; using measurements from the sensor to determine the position or orientation of the sensor with respect to the transmitter.

[0076] In some embodiments, the transmitter is integrated with a head-mounted portion of the virtual or augmented reality display system.

[0077] In some embodiments, a virtual or augmented reality display system comprises a display configured to display imagery for a plurality of depth planes; a display controller configured to receive rendered virtual or augmented reality imagery data, and to control the display based at least in part on control information embedded in the rendered imagery, wherein the embedded control information indicates a desired brightness or color to apply to at least a portion of the rendered imagery when displaying the imagery. The desired brightness or color can alter the displayed position of one or more virtual or augmented reality objects as compared to the position of the one or more objects in the rendered imagery. The desired brightness or color can longitudinal shift at least a portion of the imagery from

one depth plane to a virtual depth plane, the virtual depth plane comprising a weighted combination of at least two depth planes.

**[0078]** In some embodiments, a virtual or augmented reality display system comprises: a display configured to display imagery for a plurality of depth planes; a display controller configured to receive rendered virtual or augmented reality imagery data, and to control the display based at least in part on control information, wherein the control information indicates that at least one depth plane is inactive and the display controller is configured to control the display based on the indication that at least one depth plane is inactive, thereby reducing power consumption.

**[0079]** In some embodiments, the indication that at least one depth plane is inactive comprises control information comprising depth plane indicator data that specifies a plurality of active depth planes to display the imagery.

**[0080]** In some embodiments, indication that at least one depth plane is inactive comprises control information comprising depth plane indicator data that specifies that at least one depth plane is inactive.

**[0081]** In some embodiments, the control information is embedded in the rendered imagery.

**[0082]** In some embodiments, the display controller as a result of said control information that indicates that at least one depth plane is inactive causes one or more light sources to be reduced in power thereby reducing power consumption.

**[0083]** In some embodiments, a method in a virtual or augmented reality display system comprises: receiving rendered virtual or augmented reality imagery data for displaying imagery on a plurality of depth planes; receiving control information indicating that at least one depth plane is inactive; and displaying the imagery for a plurality of depth planes based at least in part on said control information indicating that at least one depth plane is inactive, thereby reducing power consumption.

**[0084]** In some embodiments, the control information comprises depth plane indicator data that specifies a plurality of active depth planes to display the imagery.

**[0085]** In some embodiments, the control information comprises depth plane indicator data that specifies at least one depth plane that is inactive.

[0086] In some embodiments, the control information is embedded in the rendered imagery.

[0087] In some embodiments, as a result of said control information indicating that at least one depth plane is inactive, one or more light sources is reduced in power thereby reducing power consumption.

[0088] In some embodiments, a virtual or augmented reality display system comprises: a display configured to display imagery for a plurality of depth planes having a plurality of color fields; a display controller configured to receive rendered virtual or augmented reality imagery data, and to control the display based at least in part on control information, wherein the control information indicates that at least one color field is inactive and the display controller is configured to control the display based on the indication that at least one color field is inactive, thereby reducing power consumption.

[0089] In some embodiments, the indication that at least one color field is inactive comprises control information comprising color field indicator data that specifies a plurality of active color fields to display the imagery.

[0090] In some embodiments, the indication that at least one color field is inactive comprises control information comprising color field indicator data that specifies that at least one color field is inactive.

[0091] In some embodiments, the control information is embedded in the rendered imagery.

[0092] In some embodiments, the display controller as a result of said control information that indicates that at least one color field is inactive causes one or more light sources to be reduced in power thereby reducing power consumption.

[0093] In some embodiments, a method in a virtual or augmented reality display system comprises: receiving rendered virtual or augmented reality imagery data for displaying imagery on a plurality of depth planes having a plurality of color fields; receiving control information indicating that at least one color field is inactive; and displaying the imagery for a plurality of color fields in a plurality of depth planes based at least in part on said control information indicating that at least one color field is inactive, thereby reducing power consumption.

[0094] In some embodiments, the control information comprises color field indicator data that specifies a plurality of active color fields to display the imagery.

[0095] In some embodiments, the control information comprises color field indicator data that specifies at least one color field that is inactive.

[0096] In some embodiments, the control information is embedded in the rendered imagery.

[0097] In some embodiments, as a result of said control information indicating that at least one color field is inactive, one or more light sources is reduced in power thereby reducing power consumption.

#### BRIEF DESCRIPTION OF THE DRAWINGS

[0098] Figure 1 illustrates a user's view of an augmented reality (AR) scene using an example AR system.

[0099] Figure 2 illustrates an example of wearable display system.

[0100] Figure 3 illustrates a conventional display system for simulating three-dimensional imagery for a user.

[0101] Figure 4 illustrates aspects of an approach for simulating three-dimensional imagery using multiple depth planes.

[0102] Figures 5A-5C illustrate relationships between radius of curvature and focal radius.

[0103] Figure 6 illustrates an example of a waveguide stack for outputting image information to a user.

[0104] Figure 7 shows an example of exit beams outputted by a waveguide.

[0105] Figure 8 illustrates an example design of a waveguide stack in which each depth plane has three associated waveguides that each output light of a different color.

[0106] Figure 9 illustrates an example timing scheme for a virtual or augmented reality system which displays light field imagery.

[0107] Figure 10 illustrates an example format for a frame of video data which includes appended control data.

[0108] Figure 11 illustrates another example format for a frame of video data which includes control data.

[0109] Figure 12 illustrates an example format for a pixel of video data which includes embedded control data.

[0110] Figure 13 illustrates how a frame of video can be separated into color components which can be displayed serially.

[0111] Figure 14 illustrates how a frame of light field video data can be separated, using depth plane indicator data, into multiple depth planes which can each be split into color components sub-frames for display.

[0112] Figure 15 illustrates an example where the depth plane indicator data of Figure 12 indicates that one or more depth planes of a frame of light field video data are inactive.

[0113] Figure 16 illustrates example drawing areas for a frame of computer-generated imagery in an augmented reality system.

[0114] Figure 17 schematically illustrates the possible motion of a user's head about two rotational axes.

[0115] Figure 18 illustrates how a user's head pose can be mapped onto a three-dimensional surface.

[0116] Figure 19 schematically illustrates various head pose regions which can be used to define gain factors for improving head pose tracking.

#### DETAILED DESCRIPTION

[0117] Virtual and augmented reality systems disclosed herein can include a display which presents computer-generated imagery to a user. In some embodiments, the display systems are wearable, which may advantageously provide a more immersive VR or AR experience. Figure 2 illustrates an example of wearable display system 80. The display system 80 includes a display 62, and various mechanical and electronic modules and systems to support the functioning of that display 62. The display 62 may be coupled to a frame 64, which is wearable by a display system user or viewer 60 and which is configured to position the display 62 in front of the eyes of the user 60. In some embodiments, a speaker 66 is

coupled to the frame 64 and positioned adjacent the ear canal of the user (in some embodiments, another speaker, not shown, is positioned adjacent the other ear canal of the user to provide for stereo/shapeable sound control). The display 62 is operatively coupled, such as by a wired or wireless connection 68, to a local data processing module 70 which may be mounted in a variety of configurations, such as fixedly attached to the frame 64, fixedly attached to a helmet or hat worn by the user, embedded in headphones, or otherwise removably attached to the user 60 (*e.g.*, in a backpack-style configuration, in a belt-coupling style configuration, etc.).

**[0118]** The local processing and data module 70 may include a processor, as well as digital memory, such as non-volatile memory (*e.g.*, flash memory), both of which may be utilized to assist in the processing and storing of data. This includes data captured from sensors, such as image capture devices (*e.g.*, cameras), microphones, inertial measurement units, accelerometers, compasses, GPS units, radio devices, and/or gyros. The sensors may be, *e.g.*, operatively coupled to the frame 64 or otherwise attached to the user 60. Alternatively, or additionally, sensor data may be acquired and/or processed using a remote processing module 72 and/or remote data repository 74, possibly for passage to the display 62 after such processing or retrieval. The local processing and data module 70 may be operatively coupled by communication links (76, 78), such as via a wired or wireless communication links, to the remote processing module 72 and remote data repository 74 such that these remote modules (72, 74) are operatively coupled to each other and available as resources to the local processing and data module 70.

**[0119]** In some embodiments, the remote processing module 72 may include one or more processors configured to analyze and process data (*e.g.*, sensor data and/or image information). In some embodiments, the remote data repository 74 may comprise a digital data storage facility, which may be available through the internet or other networking configuration in a “cloud” resource configuration. In some embodiments, all data is stored and all computations are performed in the local processing and data module, allowing fully autonomous use from a remote module.

**[0120]** In some embodiments, the computer-generated imagery provided via the display 62 can create the impression of being three-dimensional. This can be done, for

example, by presenting stereoscopic imagery to the user. In some conventional systems, such imagery can include separate images of a scene or object from slightly different perspectives. The separate images can be presented to the user's right eye and left eye, respectively, thus simulating binocular vision and its associated depth perception.

**[0121]** Figure 3 illustrates a conventional display system for simulating three-dimensional imagery for a user. Two distinct images 74 and 76, one for each eye 4 and 6, are outputted to the user. The images 74 and 76 are spaced from the eyes 4 and 6 by a distance 10 along an optical or z-axis parallel to the line of sight of the viewer. The images 74 and 76 are flat and the eyes 4 and 6 may focus on the images by assuming a single accommodated state. Such systems rely on the human visual system to combine the images 74 and 76 to provide a perception of depth for the combined image.

**[0122]** It will be appreciated, however, that the human visual system is more complicated and providing a realistic perception of depth is more challenging. For example, many viewers of conventional 3D display systems find such systems to be uncomfortable or may not perceive a sense of depth at all. Without being limited by theory, it is believed that viewers of an object may perceive the object as being "three-dimensional" due to a combination of vergence and accommodation. Vergence movements (*i.e.*, rolling movements of the pupils toward or away from each other to converge the lines of sight of the eyes to fixate upon an object) of the two eyes relative to each other are closely associated with focusing (or "accommodation") of the lenses of the eyes. Under normal conditions, changing the focus of the lenses of the eyes, or accommodating the eyes, to change focus from one object to another object at a different distance will automatically cause a matching change in vergence to the same distance, under a relationship known as the "accommodation-vergence reflex." Likewise, a change in vergence will trigger a matching change in accommodation, under normal conditions. As noted herein, many stereoscopic display systems display a scene using slightly different presentations (and, so, slightly different images) to each eye such that a three-dimensional perspective is perceived by the human visual system. Such systems are uncomfortable for many viewers, however, since they simply provide different presentations of a scene but with the eyes viewing all the image information at a single accommodated state, and thus work against the accommodation-vergence reflex. Display systems that



provide a better match between accommodation and vergence may form more realistic and comfortable simulations of three-dimensional imagery.

[0123] For example, light field imagery can be presented to the user to simulate a three-dimensional view. Light field imagery can mimic the rays of light which enter the eyes of a viewer in a real-world environment. For example, when displaying light field imagery, light rays from objects that are simulated to be perceived at a distance are made to be more collimated when entering the viewer's eyes, while light rays from objects that are simulated to be perceived nearby are made to be more divergent. Thus, the angles at which light rays from objects in a scene enter the viewer's eyes are dependent upon the simulated distance of those objects from the viewer. Light field imagery in a virtual or augmented reality system can include multiple images of a scene or object from different depth planes. The images may be different for each depth plane (*e.g.*, provide slightly different presentations of a scene or object) and may be separately focused by the viewer's eyes, thereby helping to provide the user with a comfortable perception of depth.

[0124] When these multiple depth plane images are presented to the viewer simultaneously or in quick succession, the result is interpreted by the viewer as three-dimensional imagery. When the viewer experiences this type of light field imagery, the eyes accommodate to focus the different depth planes in much the same way as they would do when experiencing a real-world scene. These focal cues can provide for a more realistic simulated three-dimensional environment.

[0125] In some configurations, at each depth plane, a full color image may be formed by overlaying component images that each have a particular component color. For example, red, green, and blue images may each be separately outputted to form each full color depth plane image. As a result, each depth plane may have multiple component color images associated with it.

[0126] Figure 4 illustrates aspects of an approach for simulating three-dimensional imagery using multiple depth planes. With reference to Figure 4A, objects at various distances from eyes 4 and 6 on the z-axis are accommodated by the eyes (4, 6) so that those objects are in focus. The eyes 4 and 6 assume particular accommodated states to bring into focus objects at different distances along the z-axis. Consequently, a particular

accommodated state may be said to be associated with a particular one of depth planes 14, such that objects or parts of objects in a particular depth plane are in focus when the eye is in the accommodated state for that depth plane. In some embodiments, three-dimensional imagery may be simulated by providing different presentations of an image for each of the eyes (4, 6), and also by providing different presentations of the image corresponding to each of the depth planes.

**[0127]** The distance between an object and the eye (4 or 6) can change the amount of divergence of light from that object, as viewed by that eye. Figures 5A-5C illustrate relationships between distance and the divergence of light rays. The distance between the object and the eye 4 is represented by, in order of decreasing distance, R1, R2, and R3. As shown in Figures 5A-5C, the light rays become more divergent as distance to the object decreases. As distance increases, the light rays become more collimated. Stated another way, it may be said that the light field produced by a point (the object or a part of the object) has a spherical wavefront curvature, which is a function of how far away the point is from the eye of the user. The curvature increases with decreasing distance between the object and the eye 4. Consequently, at different depth planes, the degree of divergence of light rays is also different, with the degree of divergence increasing with decreasing distance between depth planes and the viewer's eye 4. While only a single eye 4 is illustrated for clarity of illustration in Figures 5A-5C and other figures herein, it will be appreciated that the discussions regarding eye 4 may be applied to both eyes (4 and 6) of a viewer.

**[0128]** Without being limited by theory, it is believed that the human eye typically can interpret a finite number of depth planes to provide depth perception. Consequently, a highly believable simulation of perceived depth may be achieved by providing, to the eye, different presentations of an image corresponding to each of these limited number of depth planes.

**[0129]** Figure 6 illustrates an example of a waveguide stack for outputting image information to a user. A display system 1000 includes a stack of waveguides, or stacked waveguide assembly 178, that may be utilized to provide three-dimensional perception to the eye/brain using a plurality of waveguides (182, 184, 186, 188, 190). In some embodiments, the display system 1000 is the system 80 of Figure 2, with Figure 6 schematically showing

some parts of that system 80 in greater detail. For example, the waveguide assembly 178 may be integrated into the display 62 of Figure 2.

**[0130]** With continued reference to Figure 6, the waveguide assembly 178 may also include a plurality of features (198, 196, 194, 192) between the waveguides. In some embodiments, the features (198, 196, 194, 192) may be lenses. The waveguides (182, 184, 186, 188, 190) and/or the plurality of lenses (198, 196, 194, 192) may be configured to send image information to the eye with various levels of wavefront curvature or light ray divergence. Each waveguide level may be associated with a particular depth plane and may be configured to output image information corresponding to that depth plane. Image injection devices (200, 202, 204, 206, 208) may be utilized to inject image information into the waveguides (182, 184, 186, 188, 190), each of which may be configured, as described herein, to distribute incoming light across each respective waveguide, for output toward the eye 4. Light exits an output surface (300, 302, 304, 306, 308) of the image injection devices (200, 202, 204, 206, 208) and is injected into a corresponding input edge (382, 384, 386, 388, 390) of the waveguides (182, 184, 186, 188, 190). In some embodiments, a single beam of light (*e.g.*, a collimated beam) may be injected into each waveguide to output an entire field of cloned collimated beams that are directed toward the eye 4 at particular angles (and amounts of divergence) corresponding to the depth plane associated with a particular waveguide.

**[0131]** In some embodiments, the image injection devices (200, 202, 204, 206, 208) are discrete displays that each produce image information for injection into a corresponding waveguide (182, 184, 186, 188, 190, respectively). In some other embodiments, the image injection devices (200, 202, 204, 206, 208) are the output ends of a single multiplexed display which may, *e.g.*, pipe image information via one or more optical conduits (such as fiber optic cables) to each of the image injection devices (200, 202, 204, 206, 208).

**[0132]** A controller 210 controls the operation of the stacked waveguide assembly 178 and the image injection devices (200, 202, 204, 206, 208). In some embodiments, the controller 210 includes programming (*e.g.*, instructions in a non-transitory computer-readable medium) that regulates the timing and provision of image information to the waveguides

(182, 184, 186, 188, 190) according to, *e.g.*, any of the various schemes disclosed herein. In some embodiments, the controller may be a single integral device, or a distributed system connected by wired or wireless communication channels. The controller 210 may be part of the processing modules (70 or 72) (Figure 2) in some embodiments.

**[0133]** The waveguides (182, 184, 186, 188, 190) may be configured to propagate light within each respective waveguide by total internal reflection (TIR). The waveguides (182, 184, 186, 188, 190) may each be planar or curved, with major top and bottom surfaces and edges extending between those major top and bottom surfaces. In the illustrated configuration, the waveguides (182, 184, 186, 188, 190) may each include light redirecting elements (282, 284, 286, 288, 290) that are configured to redirect light, propagating within each respective waveguide, out of the waveguide to output image information to the eye 4. A beam of light is outputted by the waveguide at locations at which the light propagating in the waveguide strikes a light redirecting element. The light redirecting elements (282, 284, 286, 288, 290) may be reflective and/or diffractive optical features. While illustrated disposed at the bottom major surfaces of the waveguides (182, 184, 186, 188, 190) for ease of description and drawing clarity, in some embodiments, the light redirecting elements (282, 284, 286, 288, 290) may be disposed at the top and/or bottom major surfaces, and/or may be disposed directly in the volume of the waveguides (182, 184, 186, 188, 190). In some embodiments, the light redirecting elements (282, 284, 286, 288, 290) may be formed in a layer of material that is attached to a transparent substrate to form the waveguides (182, 184, 186, 188, 190). In some other embodiments, the waveguides (182, 184, 186, 188, 190) may be a monolithic piece of material and the light redirecting elements (282, 284, 286, 288, 290) may be formed on a surface and/or in the interior of that piece of material.

**[0134]** With continued reference to Figure 6, as discussed herein, each waveguide (182, 184, 186, 188, 190) is configured to output light to form an image corresponding to a particular depth plane. For example, the waveguide 182 nearest the eye may be configured to deliver collimated light, as injected into such waveguide 182, to the eye 4. The collimated light may be representative of the optical infinity focal plane. The next waveguide up 184 may be configured to send out collimated light which passes through the first lens (192; *e.g.*, a negative lens) before it can reach the eye 4; such first lens 192 may be configured to create

a slight convex wavefront curvature so that the eye/brain interprets light coming from that next waveguide up 184 as coming from a first focal plane closer inward toward the eye 4 from optical infinity. Similarly, the third up waveguide 186 passes its output light through both the first 192 and second 194 lenses before reaching the eye 4; the combined optical power of the first 192 and second 194 lenses may be configured to create another incremental amount of wavefront curvature so that the eye/brain interprets light coming from the third waveguide 186 as coming from a second focal plane that is even closer inward toward the person from optical infinity than was light from the next waveguide up 184.

[0135] The other waveguide layers (188, 190) and lenses (196, 198) are similarly configured, with the highest waveguide 190 in the stack sending its output through all of the lenses between it and the eye for an aggregate focal power representative of the closest focal plane to the person. To compensate for the stack of lenses (198, 196, 194, 192) when viewing/interpreting light coming from the world 144 on the other side of the stacked waveguide assembly 178, a compensating lens layer 180 may be disposed at the top of the stack to compensate for the aggregate power of the lens stack (198, 196, 194, 192) below. Such a configuration provides as many perceived focal planes as there are available waveguide/lens pairings. Both the light redirecting elements of the waveguides and the focusing aspects of the lenses may be static (*i.e.*, not dynamic or electro-active). In some alternative embodiments, they may be dynamic using electro-active features.

[0136] With continued reference to Figure 6, the light redirecting elements (282, 284, 286, 288, 290) may be configured to both redirect light out of their respective waveguides and to output this light with the appropriate amount of divergence or collimation for a particular depth plane associated with the waveguide. As a result, waveguides having different associated depth planes may have different configurations of light redirecting elements (282, 284, 286, 288, 290), which output light with a different amount of divergence depending on the associated depth plane. In some embodiments, as discussed herein, the light redirecting elements (282, 284, 286, 288, 290) may be volumetric or surface features, which may be configured to output light at specific angles. For example, the light redirecting elements (282, 284, 286, 288, 290) may be volume holograms, surface holograms, and/or diffraction gratings. Light redirecting elements, such as diffraction gratings, are described in

U.S. Patent Application No. 14/641,376, filed March 7, 2015, which is incorporated by reference herein in its entirety. In some embodiments, the features (198, 196, 194, 192) may not be lenses; rather, they may simply be spacers (*e.g.*, cladding layers and/or structures for forming air gaps).

**[0137]** In some embodiments, the light redirecting elements (282, 284, 286, 288, 290) are diffractive features that form a diffraction pattern, or “diffractive optical element” (also referred to herein as a “DOE”). Preferably, the DOE’s have a relatively low diffraction efficiency so that only a portion of the light of the beam is deflected away toward the eye 4 with each intersection of the DOE, while the rest continues to move through a waveguide via total internal reflection. The light carrying the image information is thus divided into a number of related exit beams that exit the waveguide at a multiplicity of locations and the result is a fairly uniform pattern of exit emission toward the eye 4 for this particular collimated beam reflecting around within a waveguide.

**[0138]** In some embodiments, one or more DOEs may be switchable between “on” states in which they actively diffract, and “off” states in which they do not significantly diffract. For instance, a switchable DOE may comprise a layer of polymer dispersed liquid crystal, in which microdroplets comprise a diffraction pattern in a host medium, and the refractive index of the microdroplets can be switched to substantially match the refractive index of the host material (in which case the pattern does not appreciably diffract incident light) or the microdroplet can be switched to an index that does not match that of the host medium (in which case the pattern actively diffracts incident light).

**[0139]** Figure 7 shows an example of exit beams outputted by a waveguide. One waveguide is illustrated, but it will be appreciated that other waveguides in the stack of waveguides 178 may function similarly. Light 400 is injected into the waveguide 182 at the input edge 382 of the waveguide 182 and propagates within the waveguide 182 by TIR. At points where the light 400 impinges on the DOE 282, a portion of the light exits the waveguide as exit beams 402. The exit beams 402 are illustrated as substantially parallel but, as discussed herein, they may also be redirected to propagate to the eye 4 at an angle (*e.g.*, forming divergent exit beams), depending on the depth plane associated with the waveguide 182. It will be appreciated that substantially parallel exit beams may be indicative of a

waveguide that corresponds to a depth plane at a large simulated distance (*e.g.*, optical infinity) from the eye 4. Other waveguides may output an exit beam pattern that is more divergent, which would require the eye 4 to accommodate to focus on a closer simulated distance and would be interpreted by the brain as light from a distance closer to the eye 4 than optical infinity.

**[0140]** Figure 8 schematically illustrates an example design of a stacked waveguide assembly in which each depth plane has three associated waveguides that each output light of a different color. A full color image may be formed at each depth plane by overlaying images in each of multiple component colors, *e.g.*, three or more component colors. In some embodiments, the component colors include red, green, and blue. In some other embodiments, other colors, including magenta, yellow, and cyan, may be used in conjunction with or may replace one of red, green, or blue. Each waveguide may be configured to output a particular component color and, consequently, each depth plane may have multiple waveguides associated with it. Each depth plane may have, *e.g.*, three waveguides associated with it: one for outputting red light, a second for outputting green light, and a third for outputting blue light.

**[0141]** With continued reference to Figure 8, depth planes 14a-14f are shown. In the illustrated embodiment, each depth plane has three component color images associated with it: a first image of a first color, G; a second image of a second color, R; and a third image of a third color, B. As a convention herein, the numbers following each of these letters indicate diopters ( $1/m$ ), or the reciprocal of the apparent distance of the depth plane from a viewer, and each box in the figures represents an individual component color image. In some embodiments, G is the color green, R is the color red, and B is the color blue. As discussed above, the perceived distance of the depth plane from the viewer may be established by the light redirecting elements (282, 284, 286, 288, 290), *e.g.* diffractive optical element (DOE), and/or by lenses (198, 196, 194, 192), which cause the light to diverge at an angle associated with the apparent distance.

**[0142]** In some arrangements, each component color image may be outputted by a different waveguide in a stack of waveguides. For example, each depth plane may have three component color images associated with it: a first waveguide to output a first color, G; a

second waveguide to output a second color, R; and a third waveguide to output a third color, B. In arrangements in which waveguides are used to output component color images, each box in the figure may be understood to represent an individual waveguide.

[0143] While the waveguides associated with each depth plane are shown adjacent to one another in this schematic drawing for ease of description, it will be appreciated that, in a physical device, the waveguides may all be arranged in a stack with one waveguide per level. Different depth planes are indicated in the figure by different numbers for diopters following the letters G, R, and B.

#### Display Timing Schemes

[0144] In some embodiments, a virtual or augmented reality system provides light field imagery by successively displaying multiple different depth planes for a given frame of video data. The system then updates to the next frame of video data and successively displays multiple different depth planes for that frame. For example, the first frame of video data can actually include three separate sub-frames of data: a far field frame D0, a midfield frame D1, and a near field frame D2. D0, D1, and D2 can be displayed in succession. Subsequently, the second frame of video data can be displayed. The second frame of video data can likewise include a far field frame, a midfield frame, and a near field frame, which are displayed successively, and so on. While this example uses three depth planes, light field imagery is not so-limited. Rather, any plural number of depth planes can be used depending, for example, upon the desired video frame rates and the capabilities of the system.

[0145] Because each frame of light field video data includes multiple sub-frames for different depth planes, systems which provide light field imagery may benefit from display panels which are capable of high refresh rates. For example, if the system displays video with a frame rate of 120 Hz but includes imagery from multiple different depth planes, then the display will need to be capable of a refresh rate greater than 120 Hz in order to accommodate the multiple depth plane images for each frame of video. In some embodiments, Liquid Crystal Over Silicon (LCOS) display panels are used, though other types of display panels can also be used (including color sequential displays and non-color sequential displays).



[0146] Figure 9 illustrates an example timing scheme for a virtual or augmented reality system which displays light field imagery. In this example, the video frame rate is 120 Hz and the light field imagery includes three depth planes. In some embodiments, the green, red, and blue components of each frame are displayed serially rather than at the same time.

[0147] A video frame rate of 120 Hz allows 8.333 ms in which to display all of the depth planes for a single frame of video. As illustrated in Figure 9, each frame of video data includes three depth planes and each depth plane includes green, red, and blue components. For example the depth plane D0 includes a green sub-frame, G0, a red sub-frame, R0, and a blue sub-frame, B0. Similarly, the depth plane D1 comprises green, red, and blue sub-frames, G1, R1, and B1, respectively, and the depth plane D2 comprises green, red, and blue sub-frames, G2, R2, and B2, respectively. Given that each video frame comprises three depth planes, and each depth plane has three color components, the allotted 8.333 ms is divided into nine segments of 0.926 ms each. As illustrated in Figure 9, the green sub-frame G0 for the first depth plane is displayed during the first time segment, the red sub-frame R0 for the first depth plane is displayed during the second time segment, and so on. The total green on-time for each frame of video is 2.778 ms. The same is true of the total red on-time and blue on-time for each video frame. It should be understood, however, that other video frame rates can also be used, in which case the specific time intervals illustrated in Figure 9 could be adjusted accordingly. While the individual color components are illustrated as having equal display times, this is not required and the ratios of the display times between the color components can be varied. Furthermore, the flashing order illustrated in Figure 9 for the depth planes and color component sub-frames is but one example. Other flashing orders can also be used. Moreover, while Figure 9 illustrates an embodiment which uses a color sequential display technology, the techniques described herein are not limited to color sequential displays.

[0148] Other display timing schemes are also possible. For example, the frame rate, number of depth planes, and color components can vary. In some embodiments, the frame rate of a virtual or augmented reality system as described herein is 80 Hz and there are three depth planes. In some embodiments, different depth planes can be displayed in different frames. For example, light field video with four depth planes can be displayed at an

effective frame rate of 60 Hz by displaying two depth planes per frame at a frame rate of 120 Hz (depth planes D0 and D1 can be displayed in the first 8.33 ms and depth planes D2 and D3 can be displayed in the next 8.33 ms – full depth information is provided in 16.7 ms, for an effective frame rate of 60 Hz). In some embodiments, the number of depth planes which are shown can vary spatially on the display. For example, a larger number of depth planes can be shown in a sub-portion of the display in the user's line of sight, and a smaller number of depth planes can be shown in sub-portions of the display located in the user's peripheral vision. In such embodiments, an eye tracker (e.g., a camera and eye tracking software) can be used to determine which portion of the display the user is looking at.

#### Control Data for Video Data

[0149] Figure 10 illustrates an example format for a frame of video data which includes appended control data. As illustrated in Figure 10, each frame of video data may comprise an array of pixel data formatted into rows and columns. In the illustrated example, there are 1280 columns and 960 rows of pixel data which form an image. Figure 10 also illustrates that control data 1010 can be appended to a frame of video data. In this example, a control packet 1010 can be appended to a frame of video data as, for example, an extra row. The first row (Row 000) comprises the control information, whereas Rows 1-960 contain the actual image. Thus, in this embodiment, the host transmits a resolution of 1280x961 to the display controller.

[0150] The display controller reads the appended control information and uses it, for example, to configure the image information 1020 sent to one or more display panels (e.g., a left-eye and a right-eye display panel). In this example, the row of control information 1010 is not sent to the display panels. Thus, while the host transmits information to the display controller with a resolution of 1280x961, the display controller removes the control information 1010 from the stream of data and transmits only the video information 1020 to the display panel(s) with a resolution of 1280x960. The imagery data can be transmitted to a display panel (e.g., an LCOS display panel) in, for example, Display Serial Interface (DSI) format. While Figure 10 illustrates that the appended control information 1010 comprises a single row appended at the beginning of each frame of video data, other amounts of control information could alternatively be appended. Further, the control

information 1010 does not necessarily have to be appended at the beginning of each frame of video data but could instead be inserted into the video data at other locations. However, appending control information at the beginning of a frame may allow the controller to more readily act on the control information at the beginning of a frame of rendered imagery prior to displaying the image data.

**[0151]** Figure 11 illustrates another example format for a frame of video data which includes control data. Figure 11 is similar to Figure 10 except that the control information 1110 is inserted in place of the first row of video data rather than being appended to the frame of video data before the first row. Thus, the first row (Row 000) of the frame comprises control information, while the remaining 959 rows comprise the actual image data 1120.

**[0152]** In this example, the host transmits information to the display controller with a resolution of 1280x960. The display controller can use the control data 1110 to configure the image information sent to the display panel(s). The display controller then transmits the frame of video data illustrated in Figure 11 to the display panel(s). However, in some embodiments, before transmitting the frame of video data illustrated in Figure 11 to the display panel(s), the display controller can remove the control information 1110 by, for example, setting that row of video data to zeros. This causes the first row of each frame of video data to appear as a dark line on the display.

**[0153]** Using the scheme illustrated in Figure 11, control information 1110 can be included with a frame of video data without changing the resolution of the information sent to the display controller. However, the trade-off in this example is that the effective display resolution is decreased due to the fact that some image data is replaced by the control data. While Figure 11 illustrates that the control data 1110 is inserted in place of the first row of image data, the control data could alternatively be inserted in place of another row in the frame.

**[0154]** The control data illustrated in, for example, Figures 10 and 11 (and later in Figure 12) can be used for a number of different purposes. For example, the control data can indicate whether a frame of video data should be displayed on the left-eye video panel or the right-eye video panel. The control data can indicate which of a plurality of depth planes the

frame of video data corresponds to. The control data can indicate the flashing order for the light field video data. For example, the control data can indicate the order in which to display each depth plane, as well as the order to display the color component sub-frames for each depth plane. In addition, there may be a need to shift pixels left/right or up/down after the content for the display has already been generated by the host. Rather than adjusting and re-rendering the image data, the control data can include pixel shift information which specifies the direction and magnitude of a pixel shift which should be carried out by the display controller.

**[0155]** Such pixel shifts can be carried out for a number of reasons. Pixel shifts can be performed in cases in which the image content needs to be moved on the display due to, for example, a user's head movement. In such cases, the content may be the same but its location within the viewing area on the display may need to be shifted. Rather than re-rendering the image content at the GPU and sending the whole set of pixels to the display controller again, the pixel shift can be applied to the image data using the pixel shift control information. As illustrated in Figures 10 and 11, the pixel shift control information can be included at the beginning of a frame. Alternatively, and/or additionally, a late update control data packet can be sent within a frame (*e.g.*, after the first row) to perform an appropriate pixel shift based on an updated head pose mid frame. This can be done using, for example, a Mobile Industry Processor Interface (MIPI) Display Serial Interface (DSI) virtual channel.

**[0156]** Pixel shifts can also be performed in cases in which the user is moving his or her head and a more accurate representation of the pixels is wanted. Rather than having the GPU re-render the image, a late shift on the display can be applied using the pixel shift approach. Any pixel shift described herein could impact a single depth plane or multiple depth planes. As already discussed herein, in some embodiments, there are differences in time between when various depth planes are displayed. During these time differences, the user may shift his or her eyes such that the viewing frustum may need to be shifted. This can be accomplished using a pixel shift for any of the depth planes.

**[0157]** The pixel shift control information can indicate a pixel shift in the X-Y direction within a frame of a single depth plane. Alternately, and/or additionally, the pixel shift control information can indicate a shift in the Z direction between depth plane buffers.

For example, an object that was previously displayed in one or more depth planes may move to another depth plane set with a Z-pixel shift. This type of shift can also include a scaler to enlarge or reduce the partial image for each depth. Assume, for example, that a displayed character is floating between two depth planes and there is no occlusion of that character with another object. Apparent movement of the character in the depth direction can be accomplished by re-drawing the character forward or backward one or more depth planes using the Z-pixel shift and scaler. This can be accomplished without re-rendering the character and sending a frame update to the display controller, resulting in a smoother motion performance at much lower computational cost.

**[0158]** The scaler can also be used to compensate for magnification effects that occur within the display as a result of, for example, the lenses 192, 194, 196, 198. Such lenses may create virtual images which are observable by the user. When a virtual object moves from one depth plane to another, the optical magnification of the virtual image can actually be opposite of what would be expected in the physical world. For example, in the physical world when an object is located at a further depth plane from the viewer, the object appears smaller than it would if located at a closer depth plane. However, when the virtual object moves from a nearer depth plan to a further depth plane in the display, the lenses may actually magnify the virtual image of the object. Thus, in some embodiments, a scaler is used to compensate for optical magnification effects in the display. A scaler can be provided for each depth plane to correct magnification effects caused by the optics. In addition, a scaler can be provided for each color if there are any scaling issues to be addressed on a per color basis.

**[0159]** In some embodiments, the maximum horizontal pixel shift can correspond to the entire panel width, while the maximum vertical pixel shift can correspond to the entire panel height. Both positive and negative shifts can be indicated by the control data. Using this pixel shift information, the display controller can shift a frame of video data left or right, up or down, and forward or backward between depth planes. The pixel shift information can also cause a frame of video data to be completely or partially shifted from the left-eye display panel to the right-eye display panel, or vice versa. Pixel shift information can be included for each of the depth planes in the light field video data.

**[0160]** In some embodiments, such as those where scanning-based displays are used, incremental distributed pixel shifts can be provided. For example, the images for a frame of video can be shifted incrementally in one or more depth planes until reaching the end (*e.g.*, bottom) of the image. The pixels which are displayed first can be shifted more or less than later-displayed pixels within a frame in order to compensate for head movement or in order to simulate motion of the object. Further, there can be an incremental pixel shift on a per-plane basis. For example, pixels in one depth plane can be shifted more or less than pixels in another depth plane. In some embodiments, eye tracking technology is used to determine which portion of a display screen the user is fixated on. Objects in different depth planes, or even at different locations within a single depth plane, can be pixel shifted (or not shifted) depending on where the user is looking. If there are objects that the user is not fixating on, pixel shift information for those objects may be disregarded in order to improve performance for pixel shifts in the imagery that the user is fixating on. Again, an eye tracker can be used to determine where on the display the user is looking.

**[0161]** The control information can also be used to specify and/or regulate one or more virtual depth planes. A virtual depth plane can be provided at a desired interval between two defined depth planes in a virtual or augmented reality system by blending the two depth plane images with appropriate weightings to maintain the desired brightness of the imagery. For example, if a virtual depth plane is desired between depth plane D0 and depth plane D1, then a blending unit can weight the pixel values of the D0 image data by 50% while also weighting the pixel values of the D1 image data by 50%. (So long as the weightings sum to 100%, then the apparent brightness of the imagery can be maintained.) The result would be a virtual depth plane that appears to be located midway between D0 and D1. The apparent depth of the virtual depth plane can be controlled by using different blending weights. For example, if it is desired that the virtual depth plane appear closer to D1 than D0, then the D1 image can be weighted more heavily. One or more scalers can be used to ensure that a virtual object is substantially the same size in both of the depth planes that are being blended so that like portions of the virtual object are combined during the blending operation. The control data can specify when virtual depth plane imagery is to be calculated and the control information can also include blending weights for the virtual depth

planes. In various embodiments, the weights can be stored in a programmable look up table (LUT). The control information can be used to select the appropriate weights from the LUT that would provide a desired virtual depth plane.

**[0162]** The control information can also indicate whether an image frame for one of two stereo displays should be copied into the other. For example, in the case of the most distant simulated depth plane (*e.g.*, background imagery), there may be relatively little difference (*e.g.*, due to parallax shift) between the right and left eye images. In such cases, the control information can indicate that the imagery for one of the stereo displays be copied to the other display for one or more depth planes. This can be accomplished without re-rendering the image data at the GPU for both the right and left eye displays or re-transferring data to the display controller. If there are relatively small differences between the right and left eye images, pixel shifts can also be used to compensate without re-rendering or re-transferring image data for both eyes.

**[0163]** The control data illustrated in Figures 10 and 11 can also be used for other purposes besides those specifically enumerated here.

**[0164]** While Figures 10 and 11 illustrate that rows of control data can be included with video data, control data can also (or alternatively) be embedded in individual pixels of video data. This is illustrated in Figure 12, which illustrates an example format for a pixel 1200 of video data which includes embedded control data 1240. Figure 12 illustrates that the pixel of video data comprises a blue value 1230 (Byte 0), a green value 1220 (Byte 1), and a red value 1210 (Byte 2). In this embodiment, each of the color values has a color depth of eight bits. In some embodiments, one or more of the bits corresponding to one or more of the color values can be replaced by control data 1240 at the expense of the bit depth of the color value(s). Thus, control data can be embedded directly in pixels of video data at the expense of dynamic range of the color value(s) for the pixel. For example, as illustrated in Figure 12, the highlighted two least significant bits of the blue value can be dedicated as control data 1240. Though not illustrated, bits of the other color values can also be dedicated as control data. Moreover, different numbers of pixel bits can be dedicated as control data.

**[0165]** In some embodiments, the control data 1240 embedded in the pixels of video data can be depth plane indicator data (though the control data embedded in the pixels

can also be any other type of control data, including other types described herein). As discussed herein, light field video data can include a number of depth planes. The bit depth for one or more pixels in the video frame can be reduced and the resulting available bit(s) can be used to indicate the depth plane to which a pixel corresponds.

**[0166]** As a concrete example, consider the 24-bit RGB pixel data illustrated in Figure 12. Each of the red, green, and blue color values has a bit depth of eight bits. As already discussed, the bit depth of one or more of the color components can be sacrificed and replaced by depth plane indicator data. For example, since the eye is less sensitive to blue, the blue component can be represented by six bits (bits B3-B8 in Figure 12) instead of eight. The resulting extra two bits (bits B1 and B2) can be used to specify which of up to four depth planes that pixel corresponds to. If there are more or fewer depth planes, then a greater or lesser number of color bits can be sacrificed. For example if the bit depth is reduced by one bit, up to two depth planes can be specified. If the bit depth is reduced by three bits, up to eight depth planes can be specified, etc. In this way, the dynamic range of a color value can be traded off for the ability to encode depth plane indicator data directly within the imagery data itself.

**[0167]** In some embodiments, depth plane indicator data 1240 is encoded in every pixel of video data. In other embodiments, depth plane indicator data 1240 may be encoded in one pixel per frame, or one pixel per line, one pixel per virtual or augmented reality object, etc. In addition, depth plane indicator data 1240 can be encoded in just a single color component, or in multiple color components. Similarly, the technique of encoding depth plane indicator data 1240 directly within imagery data is not limited solely to color imagery. The technique can be practiced in the same way for grayscale images, etc.

**[0168]** Figure 12 illustrates one technique for encoding depth plane indicator data in image data. Another technique is to employ chroma subsampling and use the resulting available bits as depth plane indicator data. For example, the image data can be represented in YCbCr format, where Y represents the luminance component (which may or may not be gamma corrected), Cb represents a blue-difference chroma component, and Cr represents a red-difference chroma component. Since the eye is less sensitive to chroma resolution than luminance resolution, the chroma information can be provided with a lesser resolution than



the luminance information without unduly degrading image quality. In some embodiments, a YCbCr 4:2:2 format is used in which a Y value is provided for each pixel but Cb and Cr values are each only provided for every other pixel in alternating fashion. If a pixel (in the absence of chroma subsampling) normally consists of 24 bits of information (8-bit Y value, 8-bit Cb value, and 8-bit Cr value), then after employing chroma subsampling each pixel will only require 16 bits of information (8-bit Y value and 8-bit Cb or Cr value). The remaining 8 bits can be used as depth plane indicator data. The depth plane indicator data can be used to separate the pixels into the appropriate depth planes to be displayed at the appropriate times.

**[0169]** In both the embodiment illustrated in Figure 12 and the chroma subsampling embodiment, the depth plane indicator data can specify actual depth planes supported by the virtual or augmented reality system and/or virtual depth planes as discussed herein. If the depth plane indicator data specifies a virtual depth plane, it can also specify the weightings of the depth planes to be blended, as discussed herein.

**[0170]** The usage of the embedded depth plane indicator data in the display controller is illustrated in Figure 14. But first, Figure 13 is provided by way of background to show the operation of the display controller when only a single depth plane is present. Figure 13 illustrates how a frame of video can be separated into color components which can be displayed serially. The left-hand panel 1310 of Figure 13 shows an image which comprises one frame of a 120 frame per second video. As indicated by the right-hand panel 1330 of Figure 13, the image is separated into red, green, and blue color components which are flashed on the display by the display controller over the course of 1/120 of a second (8.33 ms). For simplicity, Figure 13 shows that each of the color components is flashed once and that each of the color components is active for the same amount of time. The human vision system then fuses the individual color component sub-frames into the original color image shown in the left-hand panel of Figure 13. Figure 14 shows how this process can be adapted when each frame of video data includes multiple depth planes.

**[0171]** Figure 14 illustrates how a frame of light field video data can be separated, using depth plane indicator data, into multiple depth planes which can each be split into color components sub-frames for display. In some embodiments, a host transmits a stream of light field video data to a display controller. This stream of video data is represented by the image

in the left-hand panel 1410 of Figure 14. The display controller can use embedded depth plane indicator data 1240 to separate the stream of video data into a plurality of RxGxBx sequences, where a R0G0B0 sequence corresponds to a first depth plane, a R1G1B1 sequence corresponds to a second depth plane, and a R2G2B2 sequence corresponds to a third depth plane. As illustrated in Figure 13, this depth plane separation can be performed on the basis of the two least significant blue bits in each pixel. The result is shown in the middle panel 1420 of Figure 14, which shows three separate depth plane images. Finally, each of the three separate depth plane images shown in the middle panel 1420 of Figure 14 can be separated into its constituent color component sub-frames. The color component sub-frames of each depth plane can then be sequentially flashed to the display, as illustrated by the right-hand panel 1430 of Figure 14. The sequence order can be, for example, R0G0B0R1G1B1R2G2B2 as illustrated in Figure 14, or G0R0B0G1R1B1G2R2B2 as illustrated in Figure 9.

**[0172]** The depth plane indicator data 1240 can be used by the display controller to determine the number of RxGxBx sequences to use and which pixels correspond to which sequence. Control data can also be provided to specify the order of RxGxBx color sequences that are flashed to the display. For example, in the case of video data which includes three depth planes (D0, D1, D2), there are six possible orders in which the individual RxGxBx sequences can be flashed to the display panel: D0, D1, D2; D0, D2, D1; D1, D0, D2; D1, D2, D0; D2, D0, D1; and D2, D1, D0. If the order specified by the control data is D0, D1, D2, then pixels with blue LSB bits 0b00 corresponding to the first depth plane, D0, can be selected as the first RxGxBx color sequence image going out. Pixels with blue LSB bits 0b01 corresponding to the second depth plane, D1, can be selected as the second RxGxBx color sequence image going out, and so on.

**[0173]** Figure 15 illustrates an example where the depth plane indicator data of Figure 12 indicates that one or more depth planes of a frame of light field video data are inactive. Figure 15 is similar to Figure 14 in that it shows a stream of video data (represented by the left-hand panel 1510 of Figure 15) being separated into depth planes (represented by the middle panel 1520 of Figure 15), which are then each separated into color component

sub-frames (represented by the right-hand panel 1530 of Figure 15). However, Figure 15 is distinct from Figure 14 in that only a single depth plane is shown as being active.

[0174] As already discussed, the depth plane indicator data 1240 in Figure 12 comprises the two least significant bits of the blue value in each pixel. These two bits are capable of specifying up to four depth planes. However, light field video data may include fewer than four depth planes. For instance, in the preceding example, the light field video data includes only three depth planes. In such cases where the video data includes fewer than the maximum number of specifiable depth planes, the depth plane indicator data can specify that one or more depth planes are inactive. For example, continuing with the preceding example, if the two blue LSB bits in a pixel are set to 0b11, then the pixel can be assigned to an inactive fourth depth plane D3. As shown in Figure 15, only one of three RxGxBx color sequences is activated in the output sequence; the inactive depth planes are shown as black frames. As before, control data can be provided to specify the order in which depth planes are displayed. As shown in the middle panel 1520 of Figure 15, in the illustrated example, the control data has specified that the inactive depth plane D3 be shown first and last in the sequence. Thus, only the middle frame in the sequence comprises actual image data which is flashed to the display. (Other sequences can also be used. For example, the active depth plane could be ordered first or last in the sequence, or it could be repeated in the sequence more than once.) When the display controller sees that a pixel is assigned to an inactive depth plane, then the display controller can simply disregard the pixel and not flash it to the display. For example, when the control data indicates that a depth plane is inactive, power to the light source(s) that provides light to the display for that particular depth plane can be reduced (e.g., shut off), thereby reducing power consumption. This can save switching power at the display driver. Thus, a power-saving mode can be implemented by designating one or more depth planes of the video data as inactive. Likewise, in some embodiments, the control data can indicate that one or more color fields is inactive within a depth plane, while one or more other color fields in a depth plane are active. Based on this control data, the display controller can control the display to disregard the color field or fields that are inactive and display the imagery from the one or more active color fields without the inactive color field(s). For example, when the control data indicates that a color field is inactive, power to

the light source(s) that provides light to the display for that particular color field can be reduced (e.g., shut off), thereby reducing power consumption. Accordingly, light sources, such as light emitting diodes (LEDs), lasers, etc., that provide illumination to the display can be shut off or have their power reduced for inactive depth planes and/or inactive color fields.

#### Multi-Depth Plane Image Compression

[0175] In some embodiments, image compression techniques are applied across multiple depth planes in order to reduce the amount of video image data by removing redundancy of information between depth planes. For example, rather than transmitting an entire frame of image data for each depth plane, some or all of the depth planes may instead be represented in terms of changes with respect to an adjacent depth plane. (This can also be done on a temporal basis between frames at adjacent instants in time.) The compression technique can be lossless or it can be lossy, such that changes between adjacent depth plane frames, or between temporally-adjacent frames, which are less than a given threshold can be ignored, thus resulting in a reduction in image information. In addition, the compression algorithms can encode motion of objects within a single depth plane (X-Y motion) and/or between depth planes (Z motion) using motion vectors. Rather than requiring that image data for a moving object be repeatedly transmitted over time, motion of the object can be achieved entirely or partially with pixel shift control information, as discussed herein.

#### Dynamically Configurable Display Drawing Areas

[0176] In systems that display light field imagery, it can be challenging to achieve high video frame rates owing to the relatively large amount of information (e.g., multiple depth planes, each with multiple color components) included for each video frame. However, video frame rates can be improved, particularly in augmented reality mode, by recognizing that computer-generated light field imagery may only occupy a fraction of the display at a time, as shown in Figure 16.

[0177] Figure 16 illustrates example drawing areas for a frame of computer-generated imagery in an augmented reality system. Figure 16 is similar to Figure 1 except that it shows only the portions of the display where augmented reality imagery is to be drawn. In this case, the augmented reality imagery includes the robot statue 1110 and the bumblebee character 2. The remaining area of the display in augmented reality mode may simply be a

view of the real-world environment surrounding the user. As such, there may be no need to draw computer-generated imagery in those areas of the display. It may often be the case that the computer-generated imagery occupies only a relatively small fraction of the display area at a time. By dynamically re-configuring the specific drawing area(s) which are refreshed from frame-to-frame so as to exclude areas where no computer-generated imagery need be shown, video frame rates can be improved.

**[0178]** Computer-generated augmented reality imagery may be represented as a plurality of pixels, each having, for example, an associated brightness and color. A frame of video data may comprise an  $m \times n$  array of such pixels, where  $m$  represents a number of rows and  $n$  represents a number of columns. In some embodiments, the display of an augmented reality system is at least partially transparent so as to be capable of providing a view of the user's real-world surroundings in addition to showing the computer-generated imagery. If the brightness of a given pixel in the computer-generated imagery is set to zero or a relatively low value, then the viewer will see the real-world environment at that pixel location. Alternatively, if the brightness of a given pixel is set to a higher value, then the viewer will see computer-generated imagery at that pixel location. For any given frame of augmented reality imagery, the brightness of many of the pixels may fall below a specified threshold such that they need not be shown on the display. Rather than refresh the display for each of these below-threshold pixels, the display can be dynamically configured not to refresh those pixels.

**[0179]** In some embodiments, the augmented reality system includes a display controller for controlling the display. The controller can dynamically configure the drawing area for the display. For example, the controller can dynamically configure which of the pixels in a frame of video data are refreshed during any given refresh cycle. In some embodiments, the controller can receive computer-generated imagery data corresponding to a first frame of video. As discussed herein, the computer-generated imagery may include several depth planes. Based on the imagery data for the first frame of video, the controller can dynamically determine which of the display pixels to refresh for each of the depth planes. If, for example, the display utilizes a scanning-type display technology, the controller can dynamically adjust the scanning pattern so as to skip areas where the augmented reality

imagery need not be refreshed (*e.g.*, areas of the frame where there is no augmented reality imagery or the brightness of the augmented reality imagery falls below a specified threshold).

**[0180]** In this way, based upon each frame of video data that is received, the controller can identify a sub-portion of the display where augmented reality imagery should be shown. Each such sub-portion may include a single contiguous area or multiple non-contiguous areas (as shown in Figure 16) on the display. Such sub-portions of the display can be determined for each of the depth planes in the light field imagery data. The display controller can then cause the display to only refresh the identified sub-portion(s) of the display for that particular frame of video. This process can be performed for each frame of video. In some embodiments, the controller dynamically adjusts the areas of the display which will be refreshed at the beginning of each frame of video data.

**[0181]** If the controller determines that the area of the display which should be refreshed is becoming smaller over time, then the controller may increase the video frame rate because less time will be needed to draw each frame of augmented reality data. Alternatively, if the controller determines that the area of the display which should be refreshed is becoming larger over time, then it can decrease the video frame rate to allow sufficient time to draw each frame of augmented reality data. The change in the video frame rate may be inversely proportional to the fraction of the display that needs to be filled with imagery. For example, the controller can increase the frame rate by 10 times if only one tenth of the display needs to be filled.

**[0182]** Such video frame rate adjustments can be performed on a frame-by-frame basis. Alternatively, such video frame rate adjustments can be performed at specified time intervals or when the size of the sub-portion of the display to be refreshed increases or decreases by a specified amount. In some cases, depending upon the particular display technology, the controller may also adjust the resolution of the augmented reality imagery shown on the display. For example, if the size of the augmented reality imagery on the display is relatively small, then the controller can cause the imagery to be displayed with increased resolution. Conversely, if the size of the augmented reality imagery on the display is relatively large, then the controller can cause imagery to be displayed with decreased resolution.

### Enhanced Head Pose Estimation

[0183] As discussed herein, virtual and augmented reality systems can include body-mounted displays, such as a helmet, glasses, goggles, etc. In addition, virtual augmented reality systems can include sensors such as gyroscopes, accelerometers, etc. which perform measurements that can be used to estimate and track the position, orientation, velocity, and/or acceleration of the user's head in three dimensions. The sensors can be provided in an inertial measurement unit worn by the user on his or her head. In this way, the user's head pose can be estimated. Head pose estimates can be used as a means of allowing the user to interact with the virtual or augmented reality scene. For example, if the user turns or tilts his or her head, then the virtual or augmented reality scene can be adjusted in a corresponding manner (*e.g.*, the field of view of the scene can be shifted or tilted).

[0184] Figure 17 schematically illustrates the possible motion of a user's head about two rotational axes. As illustrated, the user can rotate his or her head about a vertical axis and a horizontal axis perpendicular to the page. Though not illustrated, the user can also rotate his or her head about a horizontal axis that lies in the plane of the page. In some embodiments, it may be useful to define the direction of the user's line of sight as the head pose direction. (Although such a definition of head pose would not necessarily account for the side tilt of the head, other definitions of head pose could.) Figure 18 illustrates how a user's head pose can be mapped onto a three-dimensional surface 1810. Figure 18 includes a surface normal vector 1820 which indicates the user's head pose. Each possible surface normal vector 1820 on the three-dimensional surface corresponds to a distinct head pose. In Figure 18, a surface normal vector pointing directly up would correspond to the user's neutral head pose when he or she is looking directly forward.

[0185] Various algorithms can be used to estimate and track the user's head pose based on the sensor measurements from the head-mounted inertial measurement unit. These include, for example, Kalman filters and other similar algorithms. These types of algorithms typically produce estimates which are based on sensor measurements over time rather than solely at any single instant. A Kalman filter, for example, includes a prediction phase where the filter outputs a predicted estimate of the head pose based on the head pose estimate at the previous instant. Next, during an update phase, the filter updates the head pose estimate

based on current sensor measurements. Such algorithms can improve the accuracy of head pose estimates, which reduces error in displaying virtual or augmented reality imagery appropriately in response to head movements. Accurate head pose estimates can also reduce latency in the system.

**[0186]** Typically, a Kalman filter or similar algorithm produces the most accurate head pose estimates for head poses near the user's neutral head pose (corresponding to a vertical surface normal vector 1820 in Figure 18). Unfortunately, such algorithms may fail to properly estimate head pose movement as the head pose deviates further from the neutral head pose because they do not account for movement limits imposed by human physiology or the movement of the user's head in relation to the body. However, various adaptations can be made in order to reduce the effects of these weaknesses on head pose tracking.

**[0187]** In some embodiments, head pose estimation and tracking using Kalman filters or similar algorithms can be improved by using variable gain factors which are different depending upon the current head pose location within an envelope of physiologically-possible head poses. Figure 18 illustrates a three-dimensional surface 1810 corresponding to such an envelope of physiologically-possible head poses. Figure 18 shows that the user's head has a range of motion in any direction of no more than about 180° (*e.g.*, side to side or up and down). The current head pose within the physiological envelope can be used to adjust the Kalman filter estimated variable gain factors. In areas near the center of the envelope (*i.e.*, neutral head pose), the gain factors can be set to emphasize the predicted head pose over the measured head pose because the Kalman filter prediction errors can be lower due to the higher linearity of the head movement in this region. This can reduce latency in the system without unduly impacting head pose estimation accuracy. When the head pose approaches the physiological head movement envelope boundaries, then the algorithm can use gain factors which are set to reduce the filter's reliance on predicted head pose or emphasize the measured head pose over the predicted head pose in order to reduce error.

**[0188]** In some embodiments, each location on the physiological head pose envelope illustrated in Figure 18 can correspond to different gains. In other embodiments,



the physiological head pose envelope can be split into separate regions and different gain values can be associated with each of the different regions. This is illustrated in Figure 19

[0189] Figure 19 schematically illustrates various head pose regions which can be used to define gain factors for improving head pose tracking. Figure 19 shows a central region 1910 corresponding to relatively neutral head poses. It also includes an outer region 1930 corresponding to head poses near the physiological boundary and an intermediate region 1920 in between the central and outer regions. In some embodiments, a different set of gain factors can be specified for each head pose region. The central region 1910 shows the areas with the higher linearity of movement which will have higher accuracy prediction values produced by a Kalman filter algorithm. When the head pose is within the central region 1910, the gain factors of the Kalman filter can be set to emphasize the predicted head pose over the measured head pose or to otherwise reduce reliance on measured head pose. As the head pose exits the central region and enters the intermediate or outer regions (1920, 1930, respectively), the movement can become more constrained by physiological factors that will adversely impact the Kalman predicted head pose if not taken into account by the algorithm. Accordingly, in these regions (particularly the outer region 1930), the Kalman filter gain values can be set to reduce the filter's reliance on predicted head pose and increase its reliance on measured head pose. For example, it would be inaccurate to strongly rely on a predicted head pose too far into the future if it is known that the acceleration of the head will come to a stop close to the envelope boundaries. Although three head pose regions are illustrated in Figure 19, a different number of head pose regions can be used in other embodiments.

[0190] In some embodiments, head pose estimation and tracking can also be improved by sensing the position, orientation, velocity, and/or acceleration of the user's head relative to the user's body rather than sensing the movement of the head in an absolute sense. This can be done by providing an additional inertial measurement unit worn by the user on his or her body (*e.g.*, on the torso or waist). It is important to note that head pose is a function of both head and body movement. The envelope of physiologically-possible head poses is not fixed in space; it moves with, for example, body rotation. If the user were sitting in a chair moving his or her head while keeping the body immobilized, then the physiological

envelope would be relatively constrained such that relatively good head pose estimates could be achieved by considering only the head movement. However, when a user is actually wearing a virtual or augmented reality head-mounted display and moving around, then the physiological envelope of possible head poses varies with body movement.

**[0191]** A second inertial measurement unit worn on the body (*e.g.*, mounted with the battery pack and/or processor for the virtual or augmented reality system) can help provide additional information to track the movement of the physiological envelope of head poses. Instead of fixing the envelope in space, the second inertial measurement unit can allow for movement of the head to be determined in relation to the body. For example, if the body rotates to the right, then the physiological envelope can be correspondingly rotated to the right to more accurately determine the head pose within the physiological envelope and avoid unduly constraining the operation of the Kalman filter.

**[0192]** In some embodiments, the motion of the head determined using the head-mounted inertial measurement unit can be subtracted from the motion of the body determined using the body-mounted inertial measurement unit. For example, the absolute position, orientation, velocity, and/or acceleration of the body can be subtracted from the absolute position, orientation, velocity, and/or acceleration of the head in order to estimate the position, orientation, velocity, and/or acceleration of the head in relation to the body. Once the orientation or motion of the head in relation to the body is known, then the actual head pose location within the physiological envelope can be more accurately estimated. As discussed herein, this allows Kalman filter gain factors to be determined in order to improve estimation and tracking of the head pose.

#### Enhanced “Totem” Position Estimation

**[0193]** In some virtual or augmented reality systems, a specified tangible object can be used as a “totem” which allows a user to interact with a virtual object or scene. For example, a tangible block which the user holds in his or her hand could be recognized by the system as an interactive device, such as a computer mouse. The system can include, for example, a camera which tracks the movement of the tangible block in the user’s hand and then accordingly adjusts a virtual pointer. A possible drawback of using computer vision for tracking totems in space is that the totems may occasionally be outside the field of view of

the camera or otherwise obscured. Thus, it would be beneficial to provide a system for robustly tracking the position and motion of the totem in three dimensions with six degrees of freedom.

[0194] In some embodiments, a system for tracking the position and motion of the totem includes one or more sensors in the totem. These one or more sensors could be accelerometers and/or gyroscopes which independently determine the position and movement of the totem in space. This data can then be transmitted to the virtual or augmented reality system.

[0195] Alternatively, the one or more sensors in the totem can work in conjunction with a transmitter to determine the position and movement of the totem and space. For example, the transmitter can create spatially-varying electric and/or magnetic fields in space and the totem can include one or more sensors which repeatedly measure the field at the location of the totem, thereby allowing the position and motion of the totem to be determined. In some embodiments, such a transmitter can advantageously be incorporated into the head-mounted display of the virtual or augmented reality system. Alternatively, the transmitter could be incorporated into a body-mounted pack. In this way, the location and/or movement of the totem with respect to the head or body, respectively, of the user can be determined. This may be more useful information than if the transmitter were simply located at a fixed location (*e.g.*, on a nearby table) because the location and/or movement of the totem can be determined in relation to the head or body of the user.

#### Adjustment of Imagery Colors Based on Ambient Lighting

[0196] In some embodiments, the virtual and augmented reality systems described herein include one or more sensors (*e.g.*, a camera) to detect the brightness and/or hue of the ambient lighting. Such sensors can be included, for example, in a display helmet of the virtual or augmented reality system. The sensed information regarding the ambient lighting can then be used to adjust the brightness or hue of generated pixels for virtual objects. For example, if the ambient lighting has a yellowish cast, computer-generated virtual objects can be altered to have yellowish color tones which more closely match those of the real objects in the room. Such pixel adjustments can be made at the time an image is rendered by the GPU.

Alternatively, and/or additionally, such pixel adjustments can be made after rendering by using the control information discussed herein.

[0197] For purposes of summarizing the disclosure, certain aspects, advantages and features of the invention have been described herein. It is to be understood that not necessarily all such advantages may be achieved in accordance with any particular embodiment of the invention. Thus, the invention may be embodied or carried out in a manner that achieves or optimizes one advantage or group of advantages as taught herein without necessarily achieving other advantages as may be taught or suggested herein.

[0198] Embodiments have been described in connection with the accompanying drawings. However, it should be understood that the figures are not drawn to scale. Distances, angles, etc. are merely illustrative and do not necessarily bear an exact relationship to actual dimensions and layout of the devices illustrated. In addition, the foregoing embodiments have been described at a level of detail to allow one of ordinary skill in the art to make and use the devices, systems, methods, etc. described herein. A wide variety of variation is possible. Components, elements, and/or steps may be altered, added, removed, or rearranged.

[0199] The devices and methods described herein can advantageously be at least partially implemented using, for example, computer software, hardware, firmware, or any combination of software, hardware, and firmware. Software modules can comprise computer executable code, stored in a computer's memory, for performing the functions described herein. In some embodiments, computer-executable code is executed by one or more general purpose computers. However, a skilled artisan will appreciate, in light of this disclosure, that any module that can be implemented using software to be executed on a general purpose computer can also be implemented using a different combination of hardware, software, or firmware. For example, such a module can be implemented completely in hardware using a combination of integrated circuits. Alternatively or additionally, such a module can be implemented completely or partially using specialized computers designed to perform the particular functions described herein rather than by general purpose computers. In addition, where methods are described that are, or could be, at least in part carried out by computer software, it should be understood that such methods can be provided on non-transitory

computer-readable media (*e.g.*, optical disks such as CDs or DVDs, hard disk drives, flash memories, diskettes, or the like) that, when read by a computer or other processing device, cause it to carry out the method.

**[0200]** While certain embodiments have been explicitly described, other embodiments will become apparent to those of ordinary skill in the art based on this disclosure.

WHAT IS CLAIMED IS:

1. A virtual or augmented reality display system comprising:  
a display configured to display imagery for a plurality of depth planes;  
a display controller configured to receive rendered virtual or augmented reality imagery data from a graphics processor, and to control the display based at least in part on control information embedded in the rendered imagery,  
wherein the embedded control information indicates a shift to apply to at least a portion of the rendered imagery when displaying the imagery.
2. The virtual or augmented reality display system of Claim 1, wherein the shift alters the displayed position of one or more virtual or augmented reality objects as compared to the position of the one or more objects in the rendered imagery.
3. The virtual or augmented reality display system of Claim 1, wherein the shift comprises a lateral shift of at least a portion of the imagery by a specified number of pixels within the same depth plane.
4. The virtual or augmented reality display system of Claim 1, wherein the shift comprises a longitudinal shift of at least a portion of the imagery from one depth plane to another.
5. The virtual or augmented reality display system of Claim 4, wherein the display controller is further configured to scale at least a portion of the imagery in conjunction with a longitudinal shift from one depth plane to another.
6. The virtual or augmented reality display system of Claim 1, wherein the shift comprises a longitudinal shift of at least a portion of the imagery from one depth plane to a virtual depth plane, the virtual depth plane comprising a weighted combination of at least two depth planes.
7. The virtual or augmented reality display system of Claim 1, wherein the shift is based on information regarding a head pose of a user.
8. The virtual or augmented reality display system of Claim 1, wherein the shift is performed by the display controller without re-rendering the imagery.
9. A method in a virtual or augmented reality display system, the method comprising:

receiving rendered virtual or augmented reality imagery data from a graphics processor; and

displaying the imagery for a plurality of depth planes based at least in part on control information embedded in the rendered imagery,

wherein the embedded control information indicates a shift to apply to at least a portion of the rendered imagery when displaying the imagery.

10. The method of Claim 9, wherein displaying the imagery comprises shifting the displayed position of one or more virtual or augmented reality objects as compared to the position of the one or more objects in the rendered imagery.

11. The method of Claim 9, wherein shifting the displayed position comprises laterally shifting at least a portion of the imagery by a specified number of pixels within the same depth plane based on the control information.

12. The method of Claim 9, wherein shifting the displayed position comprises longitudinally shifting at least a portion of the imagery from one depth plane to another based on the control information.

13. The method of Claim 12, further comprising scaling at least a portion of the imagery in conjunction with longitudinally shifting the imagery from one depth plane to another.

14. The method of Claim 9, wherein shifting the displayed position comprises longitudinally shifting at least a portion of the imagery from one depth plane to a virtual depth plane, the virtual depth plane comprising a weighted combination of at least two depth planes.

15. The method of Claim 9, wherein the shift is based on information regarding a head pose of a user.

16. The method of Claim 9, further comprising shifting the imagery without re-rendering the imagery.

17. A virtual or augmented reality display system comprising:

a sensor configured to determine one or more characteristics of the ambient lighting;

a processor configured to adjust one or more characteristics of a virtual object based on the one or more characteristics of the ambient lighting; and  
a display configured to display the virtual object to a user.

18. The virtual or augmented reality display system of Claim 17, wherein the one or more characteristics of the ambient lighting comprise the brightness of the ambient lighting.

19. The virtual or augmented reality display system of Claim 17, wherein the one or more characteristics of the ambient lighting comprise the hue of the ambient lighting.

20. The virtual or augmented reality display system of Claim 17, wherein the one or more characteristics of the virtual object comprise the brightness of the virtual object.

21. The virtual or augmented reality display system of Claim 17, wherein the one or more characteristics of the virtual object comprise the hue of the virtual object.

22. A method in a virtual or augmented reality display system, the method comprising:

receiving one or more characteristics of the ambient lighting from a sensor;  
adjusting, using a processor, one or more characteristics of a virtual object based on the one or more characteristics of the ambient lighting; and  
displaying the virtual object to a user.






FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5A




FIG. 5B




FIG. 5C




FIG. 6




FIG. 7





FIG. 8



Frame Rate	120Hz
ms	8.333
color on time	0.926
Blue On Time	2.778
Green On Time	2.778
Red on Time	2.778
Total	8.333

FIG. 9






FIG. 10




FIG. 11

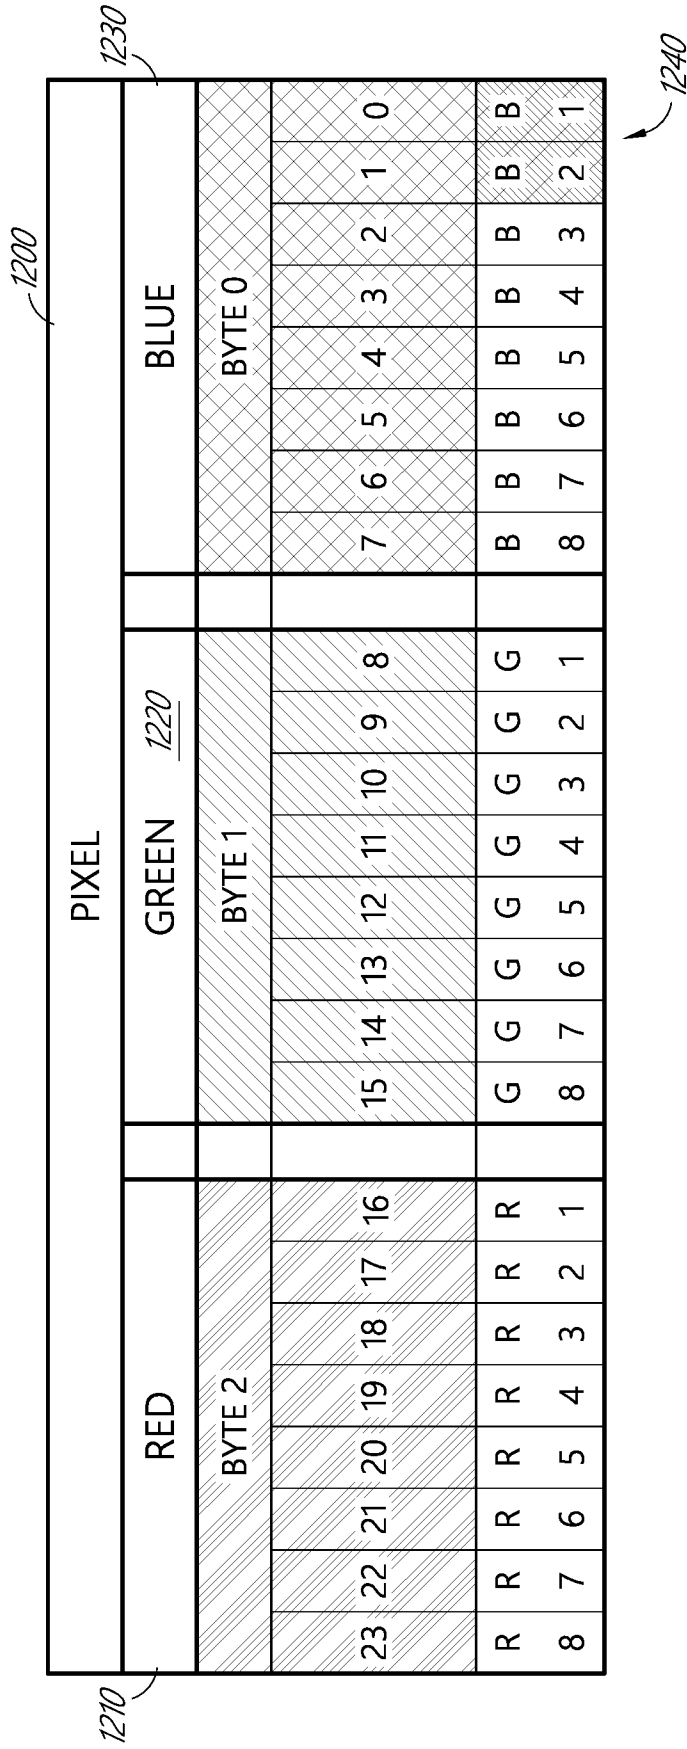


FIG. 12

12/16




FIG. 13




FIG. 14




FIG. 15




FIG. 16




FIG. 17





FIG. 18

Top view of Headform



Side view of Headform

FIG. 19