



(51) International Patent Classification:

H04N 19/167 (2014.01) H04N 19/86 (2014.01)  
H04N 19/80 (2014.01) H04N 19/63 (2014.01)

(21) International Application Number:

PCT/US2019/039595

(22) International Filing Date:

27 June 2019 (27.06.2019)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/691,423 28 June 2018 (28.06.2018) US

(71) Applicant: **APPLE INC.** [US/US]; One Apple Park Way, Cupertino, California 95014 (US).

(72) Inventors: **CHOU, Jim C.**; One Apple Park Way, Cupertino, California 95014 (US). **CISMAS, Sorin C.**; One Apple Park Way, Cupertino, California 95014 (US).

(74) Agent: **KOWERT, Robert C.**; Meyertons, Hood, Kivlin, Kowert & Goetzel, P.C., P.O. Box 398, Austin, Texas 78767-0398 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO,

DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

- with international search report (Art. 21(3))
- 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: VIDEO ENCODING SYSTEM




FIG. 1

(57) Abstract: A video encoding system in which pixel data is decomposed into frequency bands prior to encoding. The frequency bands are organized into blocks that are provided to a block-based encoder. The encoded frequency data is packetized and transmitted to a receiving device. On the receiving device, the encoded data is decoded to recover the frequency bands. Wavelet synthesis is then performed on the frequency bands to reconstruct the pixel data for display. The system may encode parts of frames (tiles or slices) using one or more encoders and transmit the encoded parts as they are ready. A pre-filter component may perform a lens warp on the pixel data prior to the wavelet transform.



## VIDEO ENCODING SYSTEM

### BACKGROUND

5

[0001] Virtual reality (VR) allows users to experience and/or interact with an immersive artificial environment, such that the user feels as if they were physically in that environment. For example, virtual reality systems may display stereoscopic scenes to users in order to create an illusion of depth, and a computer may adjust the scene content in real-time to provide the illusion of the user moving within the scene. When the user views images through a virtual reality system, the user may thus feel as if they are moving within the scenes from a first-person point of view. Similarly, mixed reality (MR) combines computer generated information (referred to as virtual content) with real world images or a real world view to augment, or add content to, a user's view of the world, or alternatively combines virtual representations of real world objects with views of a three-dimensional (3D) virtual world. The simulated environments of virtual reality and/or the mixed environments of mixed reality may thus be utilized to provide an interactive user experience for multiple applications.

10

15

### SUMMARY

[0002] Various embodiments of a video encoding system are described that may encode high-resolution video sources at low latencies for transmission over a communications link (e.g., a wireless link) to a device for decoding and display. Embodiments of the video encoding system may also provide graceful degradation of encoded video transmitted to the device to maintain a desired frame rate in varying conditions such as variations in the channel capacity of the communications link. An example application of the video encoding system is in virtual or mixed reality systems in which video frames containing virtual content are rendered, encoded, and transmitted by a base station to a device (e.g., a notebook or laptop computer, pad or tablet device, smartphone, or head-mounted display (HMD) such as a headset, helmet, goggles, or glasses that may be worn by a user) for decoding and display.

20

25

[0003] Various methods and apparatus may be implemented by the video encoding system to maintain the target frame rate through the wireless link and to minimize the latency in frame rendering, transmittal, and display. In addition, the methods and apparatus may provide graceful degradation of encoded video transmitted to the device to maintain a desired frame rate in varying conditions such as variations in the channel capacity of the communications link.

30

**[0004]** In some embodiments, the video encoding system may perform a wavelet transform on the pixel data prior to encoding to decompose the pixel data into frequency bands. The frequency bands are then organized into blocks that are provided to a block-based encoder for encoding/compression. The encoded frequency data is then sent to a wireless interface that packetizes the encoded frequency data and transmits the packets to the receiving device. On the receiving device, the encoded data is de-packetized and passed through a block-based decoder to recover the frequency bands. Wavelet synthesis is then performed on the recovered frequency bands to reconstruct the pixel data for display.

**[0005]** In some embodiments, the video encoding system may perform slice-based rendering, encoding, and transmittal. Rather than rendering and encoding entire frames in the base station and transmitting the rendered frames to the device, the base station may render and encode parts or portions of frames (referred to as slices) and transmit the encoded slices to the device as they are ready. Slice-based rendering and encoding may help to reduce latency, and may also reduce the amount of memory needed for buffering, which reduces the memory footprint on the chip(s) or processor(s) as well as power requirements.

**[0006]** In some embodiments, the video encoding system may perform tile-based rendering, encoding, and transmittal. In tile-based rendering, encoding, and transmittal, each slice may be divided into multiple tiles (e.g., four tiles), and the base station may render and encode the tiles and transmit the encoded tiles to the device as they are ready.

**[0007]** In some embodiments, the video encoding system may perform tile-based encoding using a single encoder to process tiles from each slice. However, in some embodiments, the video encoding system may perform tile-based encoding using multiple encoders to process respective tiles from each slice. For example, in some embodiments, each slice may be divided into four tiles, and two encoders may operate on two tiles from each slice. Each encoder may multiplex the processing of blocks from different frequency bands between its tiles to allow for multiple time units between the processing of blocks from the same frequency band. By multiplexing the processing of blocks between tiles, dependencies between blocks in a frequency band may be handled appropriately.

**[0008]** In some embodiments, the video encoding system may perform pre-filtering of the pixel data in frames prior to the wavelet transform. In some embodiments, pre-filtering may include performing a lens warp function on the frames on the base station prior to the wavelet transform. The lens warp is performed to correct for the distortion of the images introduced by the lenses of the device (e.g., of an HMD) that the images are viewed through, thus improving quality of the images. Performing the lens warp on the base station in the pre-filter stage may reduce the resolution of the frames prior to performing the wavelet transform and encoding, which

may help in improving compression, and may reduce latency and bandwidth usage on the wireless link. In addition, by performing the lens warp on the base station in the pre-filter stage rather than on the device after decoding, filtering of the image data may only need to be performed once, as opposed to performing filtering on the base station to reduce resolution prior to encoding and then performing lens warp filtering on the device.

[0009] In some embodiments, pre-filtering may include filtering to reduce resolution in peripheral regions while maintaining higher resolution in foveated regions. In this method, gaze tracking information obtained from the device may be used to identify the direction in which the user is currently looking. A region of the frame that corresponds to the fovea (referred to as the foveated region) may be identified based at least in part on the determined gaze direction. In some embodiments, the peripheral region (i.e. the portion of the frame outside the foveated region) may be pre-filtered to reduce resolution in the peripheral region.

### **BRIEF DESCRIPTION OF THE DRAWINGS**

[0010] FIG. 1 is a block diagram illustrating a video encoding system that decomposes pixel data into frequency bands using a wavelet transform prior to encoding, according to at least some embodiments.

[0011] FIG. 2 illustrates a video encoding system that includes multiple encoders that process tiles from frames in parallel, according to at least some embodiments.

[0012] FIG. 3A illustrates an example frame divided into slices and tiles, according to at least some embodiments.

[0013] FIG. 3B illustrates an example tile divided into blocks, according to at least some embodiments.

[0014] FIG. 3C illustrates performing a wavelet transform of a pixel block that stores pixel data to generate frequency band data prior to encoding, according to at least some embodiments.

[0015] FIG. 4 is a high-level flowchart of a method of operation for VR/MR systems that include video encoding systems as illustrated in FIGS. 1 and 2, according to at least some embodiments.

[0016] FIG. 5 is a flowchart of a method of operation for a video encoding system as illustrated in FIG. 1, according to at least some embodiments.

[0017] FIG. 6 is a flowchart of a method of operation for a video encoding system as illustrated in FIG. 2, according to at least some embodiments.

[0018] FIG. 7 illustrates an example VR/MR system that may implement a video encoding system, according to at least some embodiments.

[0019] FIG. 8 is a block diagram illustrating components of a VR/MR system as illustrated in FIG. 7, according to at least some embodiments.

[0020] This specification includes references to “one embodiment” or “an embodiment.” The appearances of the phrases “in one embodiment” or “in an embodiment” do not necessarily refer to the same embodiment. Particular features, structures, or characteristics may be combined in any suitable manner consistent with this disclosure.

[0021] “Comprising.” This term is open-ended. As used in the claims, this term does not foreclose additional structure or steps. Consider a claim that recites: “An apparatus comprising one or more processor units ...” Such a claim does not foreclose the apparatus from including additional components (e.g., a network interface unit, graphics circuitry, etc.).

[0022] “Configured To.” Various units, circuits, or other components may be described or claimed as “configured to” perform a task or tasks. In such contexts, “configured to” is used to connote structure by indicating that the units/circuits/components include structure (e.g., circuitry) that performs those task or tasks during operation. As such, the unit/circuit/component can be said to be configured to perform the task even when the specified unit/circuit/component is not currently operational (e.g., is not on). The units/circuits/components used with the “configured to” language include hardware - for example, circuits, memory storing program instructions executable to implement the operation, etc. Reciting that a unit/circuit/component is “configured to” perform one or more tasks is expressly intended not to invoke 35 U.S.C. § 112, paragraph (f), for that unit/circuit/component. Additionally, “configured to” can include generic structure (e.g., generic circuitry) that is manipulated by software or firmware (e.g., an FPGA or a general-purpose processor executing software) to operate in manner that is capable of performing the task(s) at issue. “Configure to” may also include adapting a manufacturing process (e.g., a semiconductor fabrication facility) to fabricate devices (e.g., integrated circuits) that are adapted to implement or perform one or more tasks.

[0023] “First,” “Second,” etc. As used herein, these terms are used as labels for nouns that they precede, and do not imply any type of ordering (e.g., spatial, temporal, logical, etc.). For example, a buffer circuit may be described herein as performing write operations for “first” and “second” values. The terms “first” and “second” do not necessarily imply that the first value must be written before the second value.

[0024] “Based On” or “Dependent On.” As used herein, these terms are used to describe one or more factors that affect a determination. These terms do not foreclose additional factors that may affect a determination. That is, a determination may be solely based on those factors or based, at least in part, on those factors. Consider the phrase “determine A based on B.” While in this case,

B is a factor that affects the determination of A, such a phrase does not foreclose the determination of A from also being based on C. In other instances, A may be determined based solely on B.

[0025] “Or.” When used in the claims, the term “or” is used as an inclusive or and not as an exclusive or. For example, the phrase “at least one of x, y, or z” means any one of x, y, and z, as well as any combination thereof.

### **DETAILED DESCRIPTION**

[0026] Various embodiments of a video encoding system are described. Embodiments of the video encoding system may encode high-resolution video sources at low latencies for transmission over a communications link (e.g., a wireless link) to a device for decoding and display. Embodiments of the video encoding system may also provide graceful degradation of encoded video transmitted to the device to maintain a desired frame rate in varying conditions such as variations in the channel capacity of the communications link.

[0027] An example application of the video encoding system is in virtual or mixed reality systems in which video frames containing virtual content are rendered, encoded, and transmitted to a device for decoding and display. Embodiments of a virtual or mixed reality system (referred to herein as a VR/MR system) are described in which embodiments of the video encoding system may be implemented. In some embodiments, the VR/MR system may include a device (e.g., a pad or tablet device, a smartphone, or a headset, helmet, goggles, or glasses worn by the user, referred to herein as a head-mounted display (HMD)), and a separate computing device, referred to herein as a base station. In some embodiments, the device and base station may each include wireless communications technology that allows the device and base station to communicate and exchange data via a wireless connection. In some embodiments, the device may include sensors that collect information about the user’s environment (e.g., video, depth information, lighting information, etc.) and information about the user (e.g., the user’s expressions, eye movement, hand gestures, etc.). The information collected by the sensors may be transmitted to the base station via the wireless connection. The base station may include software and hardware (e.g., processors (system on a chip (SOC), CPUs, image signal processors (ISPs), graphics processing units (GPUs), encoder/decoders (codecs), etc.), memory, etc.) configured to generate and render frames that include virtual content based at least in part on the sensor information received from the device via the wireless connection. The base station may also include an embodiment of the video encoding system as described herein that may pre-filter, compress and transmit the rendered frames to the device for display via the wireless connection.

**[0028]** In some embodiments, the VR/MR system may implement a proprietary wireless communications technology that provides a highly directional wireless link between the device and the base station. In some embodiments, the directionality and bandwidth of the wireless communication technology may support multiple devices communicating with the base station at the same time to thus enable multiple users to use the system at the same time in a co-located environment. However, other commercial (e.g., Wi-Fi, Bluetooth, etc.) or proprietary wireless communications technologies may be supported in some embodiments.

**[0029]** Primary constraints to be considered on a wireless link include bandwidth and latency. A target of the VR/MR system is to provide a high resolution, wide field of view (FOV) virtual display at a frame rate to provide the user with a high-quality VR/MR view. Another target is to minimize latency between the time a frame is rendered by the base station and the time the frame is displayed by the device. However, the channel capacity of the wireless link may vary with time, and the wireless link may thus support only a certain amount of information to be transmitted at any given time.

**[0030]** Various methods and apparatus are described herein that may be implemented by the video encoding system to maintain the target frame rate through the wireless link and to minimize the latency in frame rendering, transmittal, and display. In addition, the methods and apparatus may provide graceful degradation of encoded video transmitted to the device to maintain a desired frame rate in varying conditions such as variations in the channel capacity of the communications link.

**[0031]** In some embodiments, the video encoding system may perform a wavelet transform on the pixel data prior to encoding to decompose the pixel data into frequency bands. The frequency bands are then organized into blocks that are provided to a block-based encoder for encoding/compression. As an example, a frame may be divided into 128x128 blocks, and a two-level wavelet decomposition may be applied to each 128x128 block to generate 16 32x32 blocks of frequency data representing seven frequency bands that may then be sent to an encoder (e.g., a High Efficiency Video Coding (HEVC) encoder) to be encoded. The encoded frequency data is then sent to a wireless interface that packetizes the encoded frequency data and transmits the packets to the receiving device (e.g., an HMD). On the receiving device, the encoded data is de-packetized and passed through a block-based decoder to recover the frequency bands. Wavelet synthesis is then performed on the recovered frequency bands to reconstruct the pixel data for display.

**[0032]** While embodiments are generally described in which the wavelet transform is a two-level wavelet decomposition applied to each pixel block from a video frame, in various embodiments the wavelet decomposition may be any number N of levels (e.g., one level, two

levels, three levels, four levels, etc.), and the number N of levels may be adjusted to trade-off quality of the encoded image vs. complexity of the blocks to be encoded.

**[0033]** In some embodiments, a coring function may be applied to each of the wavelet coefficients within the wavelet decomposition. “Coring” refers to soft-thresholding of the coefficients, and is effective at removing noise. This may be done in an adaptive manner depending upon regions of importance vs. available bandwidth.

**[0034]** In some embodiments, the video encoding system may perform slice-based rendering, encoding, and transmittal. Rendering, encoding, and transmitting entire frames may have a latency and memory impact as each frame needs to be completed, stored, and then transmitted to the next stage of the VR/MR system. In slice-based rendering, rather than rendering and encoding entire frames in the base station and transmitting the rendered frames to the device, the base station may render and encode parts of frames (referred to as slices) and transmit the encoded slices to the device as they are ready. A slice may, for example, be a row of 128x128 blocks, or two or more rows of blocks. Slice-based rendering and encoding may help to reduce latency, and may also reduce the amount of memory needed for buffering, which may reduce the memory footprint on the chip(s) or processor(s) as well as power requirements.

**[0035]** In some embodiments, the video encoding system may perform tile-based rendering, encoding, and transmittal. In tile-based rendering, encoding, and transmittal, each slice may be divided into multiple tiles (e.g., four tiles), and the base station may render and encode the tiles and transmit the encoded tiles to the device as they are ready.

**[0036]** In some embodiments, the video encoding system may perform tile-based encoding using a single encoder to process tiles from each slice. However, in some embodiments, the video encoding system may perform tile-based encoding using multiple encoders to process respective tiles from each slice. For example, in some embodiments, each slice may be divided into four tiles, each tile including multiple 128x128 blocks, and two encoders (e0 and e1) may operate on two tiles from each slice (e.g., e0 operates on t0 and t1; e1 operates on t2 and t3). Each encoder may multiplex the processing of blocks from different frequency bands between its two tiles to allow for 16 time units between the processing of blocks from the same frequency band. By multiplexing the processing of blocks between two tiles, dependencies between blocks in a frequency band may be handled appropriately. While embodiments are described in which each slice is divided into four tiles and two encoders operate on respective tiles from each slice, slices may be divided into more tiles (e.g., six or eight tiles) in some embodiments, and more encoders (e.g., three or four encoders) may be used in some embodiments.

**[0037]** In some embodiments, the video encoding system may perform pre-filtering of the pixel data in frames prior to the wavelet transform. Pre-filtering may, for example, reduce the

resolution of the frames rendered by the base station prior to performing the wavelet transform, encoding, and transmission of the frames to the device over the wireless link, which may help in improving compression, and may reduce latency and bandwidth usage on the wireless link.

**[0038]** In some embodiments in which the device is an HMD, pre-filtering may include performing a lens warp on the frames on the base station prior to the wavelet transform. The lens warp is performed to correct for the distortion of the images introduced by the lenses on the HMD that the images are viewed through, thus improving quality of the images. In some embodiments, the HMD may store lens warp data for the lenses, for example generated by a calibration process, and may provide the lens warp data to the base station over the wireless connection. The base station may then perform the lens warp on the frames based on the lens warp data for that HMD. In conventional VR/MR systems, the lens warp is performed on the HMD after decoding and prior to display. Performing the lens warp on the base station in the pre-filter stage may reduce the resolution of the frames prior to performing the wavelet transform and encoding, which may help in improving compression, and may reduce latency and bandwidth usage on the wireless link. In addition, by performing the lens warp on the base station in the pre-filter stage rather than on the HMD after decoding, filtering of the image data may only need to be performed once, as opposed to performing filtering on the base station to reduce resolution prior to encoding and then performing lens warp filtering on the HMD.

**[0039]** In some embodiments, pre-filtering may include filtering to reduce resolution in peripheral regions while maintaining higher resolution in foveated regions. In this method, gaze tracking information obtained from the device may be used to identify the direction in which the user is currently looking. Human eyes can perceive higher resolution in the foveal region than in the peripheral region. Thus, a region of the frame that corresponds to the fovea (referred to as the foveated region) may be identified based at least in part on the determined gaze direction. In some embodiments, the peripheral region (i.e. the portion of the frame outside the foveated region) may be pre-filtered to reduce information based on knowledge of the human vision system, for example by filtering high frequency information and/or increasing color compression. In some embodiments, the amount of filtering applied to the peripheral region may increase extending towards the periphery of the frame. Pre-filtering of the peripheral region may help to provide improved compression of the frame.

**[0040]** FIG. 1 is a block diagram illustrating a video encoding system 120 that decomposes pixel data into frequency bands using a wavelet transform prior to encoding, according to at least some embodiments. A VR/MR system 10 may include at least one device 150 (e.g., a pad or tablet device, a smartphone, or an HMD such as a headset, helmet, goggles, or glasses that may be worn by a user) and a computing device 100 (referred to herein as a base station). The base station 100

renders VR or MR frames including virtual content, encodes the frames, and transmits the encoded frames over a wireless connection 180 to the device 150 for decoding and display by the device 150.

**[0041]** In some embodiments, the device 150 may include sensors 160 that collect information about the user 190's environment (e.g., video, depth information, lighting information, etc.) and about the user 190 (e.g., the user's expressions, eye movement, gaze direction, hand gestures, etc.). The device 150 may transmit at least some of the information collected by sensors 160 to the base station 100 via wireless connection 180. The base station 100 may render frames for display by the device 150 that include virtual content based at least in part on the various information obtained from the sensors 160, encode the frames, and transmit the encoded frames to the device 150 for decoding and display to the user via the wireless connection 180.

**[0042]** The base station 100 and device 150 may implement wireless communications technology that allows the base station 100 and device 150 to communicate and exchange data via a wireless connection 180. In some embodiments, the wireless connection 180 may be implemented according to a proprietary wireless communications technology that provides a highly directional wireless link between the device 150 and the base station 100. However, other commercial (e.g., Wi-Fi, Bluetooth, etc.) or proprietary wireless communications technologies may be used in some embodiments.

**[0043]** Primary constraints to be considered on the wireless connection 180 between the device 150 and the base station 100 in a VR/MR system 10 include bandwidth and latency. For example, in some embodiments, a target is to provide a high resolution, wide field of view (FOV) virtual display to the user at a frame rate that provides the user with a high-quality VR/MR view. Another target is to minimize latency between the time a video frame is captured by the device and the time a rendered VR/MR frame based on the video frame is displayed by the device.

**[0044]** The base station 100 may include various hardware components for rendering, filtering, encoding, and transmitting video and/or images as described herein, for example various types of processors, integrated circuits (ICs), central processing units (CPUs), graphics processing units (GPUs), image signal processors (ISPs), encoder/decoders (codecs), etc. The base station 100 may include, but is not limited to, a GPU rendering 110 component, a wireless interface 130 component, and a video encoding system 120 that may include one or more hardware components that implement various methods that may help to maintain the target frame rate through the wireless connection 180 and to minimize the latency in frame rendering, encoding, transmittal, and display. The video encoding system 120 may include, but is not limited to, a pre-filter 122 component (e.g., an N-channel filter bank), a wavelet transform 124 component, and an encoder 126 component.

**[0045]** GPU rendering 110 may include one or more hardware components that may render frames for display by the device 150 that include virtual content based at least in part on the various information obtained from the sensors 160.

**[0046]** In some embodiments, the video encoding system 120 may include one or more hardware components that pre-filter 122 the pixel data in the rendered frames prior to performing a wavelet transform 124. Pre-filter 122 may, for example, reduce the resolution of the frames rendered on the base station 100 prior to performing the wavelet transform 124, encoding 126, and transmission to the device 150 over the wireless connection 180, which may help in improving compression, and may reduce latency and bandwidth usage on the wireless connection 180.

**[0047]** In some embodiments, pre-filter 122 may perform a lens warp on the frames on the base station 100 prior to the wavelet transform 124. The lens warp is performed to correct for the distortion of the images introduced by the lenses on the device that the images are viewed through, thus improving quality of the images. In some embodiments, the device 150 may store lens warp data for the lenses, for example generated by a calibration process, and may provide the lens warp data to the base station 100 over the wireless connection 180. The pre-filter 122 component of the video encoding system 120 may then perform the lens warp on the frames based on the lens warp data for that device 150. In conventional VR/MR systems, the lens warp is performed on the device 150 after decoding and prior to display. Performing the lens warp on the base station 100 in the pre-filter 122 stage may reduce the resolution of the frames prior to performing the wavelet transform 124 and encoding 126, which may help in improving compression, and may reduce latency and bandwidth usage on the wireless connection 180. In addition, by performing the lens warp on the base station 100 in the pre-filter 122 stage rather than on the device 150 after decoding, filtering of the image data may only need to be performed once, as opposed to performing filtering on the base station 100 to reduce resolution prior to encoding 126 and then performing lens warp filtering on the device 150.

**[0048]** In some embodiments, pre-filter 122 may also apply one or more filters to reduce resolution in peripheral regions while maintaining higher resolution in foveated regions. In this method, gaze tracking information obtained from the device 150 may be used to identify the direction in which the user is currently looking. Human eyes can perceive higher resolution in the foveal region than in the peripheral region. Thus, a region of the frame that corresponds to the fovea (referred to as the foveated region) may be identified based at least in part on the determined gaze direction. In some embodiments, the peripheral region (i.e. the portion of the frame outside the foveated region) may be pre-filtered to reduce information based on knowledge of the human vision system, for example by filtering high frequency information and/or increasing color compression. In some embodiments, the amount of filtering applied to the peripheral region may

increase extending towards the periphery of the frame. Pre-filtering of the peripheral region may help to provide improved compression of the frame.

**[0049]** In some embodiments, a wavelet transform 124 component of the video encoding system 120 may include one or more hardware components (e.g., an N-channel filter bank) that perform a wavelet transform on the pixel data prior to encoding to decompose the pixel data into frequency bands. The frequency bands are then organized into blocks that are provided to a block-based encoder 126 for encoding/compression. As an example, as illustrated in FIGS. 3A through 3C, a frame may be divided into 128x128 blocks, and a two-level wavelet decomposition may be applied to each 128x128 block to generate 16 32x32 blocks of frequency data representing seven frequency bands that may then be sent to a block-based encoder (e.g., a High Efficiency Video Coding (HEVC) encoder) 126 to be encoded. The encoded frequency data is then sent to a wireless interface 130, implemented by one or more hardware components, that packetizes the data and transmits the packets to the device 150 over a wireless connection 180.

**[0050]** The device 150 may include various hardware components for decoding and displaying video and/or images as described herein, for example various types of processors, integrated circuits (ICs), central processing units (CPUs), graphics processing units (GPUs), image signal processors (ISPs), encoder/decoders (codecs), etc. The device 150 may include, but is not limited to, a wireless interface 152, a decoder 154 component (e.g., High Efficiency Video Coding (HEVC) decoder), a wavelet synthesis 156 component, and a display 158 component. On the device 150, the wireless interface 152 receives the packets that were transmitted over the wireless connection 180 by the base station 100. The encoded data is de-packetized and passed through a block-based decoder 154 (e.g., a High Efficiency Video Coding (HEVC) decoder) to recover the frequency bands. Wavelet synthesis 156 is then performed on the recovered frequency data to reconstruct the pixel data for display 158.

**[0051]** In some embodiments, the video encoding system 120 may perform slice-based rendering, encoding, and transmittal. Rendering, encoding, and transmitting entire frames may have a latency and memory impact as each frame needs to be completed, stored, and then transmitted to the next stage of the VR/MR system 10. In slice-based rendering, rather than rendering and encoding entire frames in the base station 100 and transmitting the rendered frames to the device 150, the base station 100 may render and encode parts of frames (referred to as slices) and transmit the encoded slices to the device 150 as they are ready. A slice may, for example, be a row of 128x128 blocks. Slice-based rendering and encoding may help to reduce latency, and may also reduce the amount of memory needed for buffering, which reduces the memory footprint on the chip(s) or processor(s) as well as power requirements.

**[0052]** In some embodiments, the video encoding system 120 may perform tile-based rendering, encoding, and transmittal. In tile-based rendering, encoding, and transmittal, each slice may be divided into multiple tiles (e.g., four tiles), and the base station 100 may render and encode the tiles and transmit the encoded tiles to the device 150 as they are ready.

5 **[0053]** In some embodiments, the video encoding system 120 may perform tile-based rendering, encoding, and transmittal using a single encoder 126 to process tiles from each slice. However, in some embodiments, the video encoding system 120 may perform tile-based encoding using multiple encoders 126 to process respective tiles from each slice. FIG. 2 illustrates a video encoding system 220 that includes multiple encoders (two encoders 226A and 226B, in this  
10 example) that process tiles from rendered frames in parallel, according to at least some embodiments.

**[0054]** A GPU rendering 210 component of the base station 200 may include one or more GPUs and/or other components that render frames (or slices of frames) for display. A frame may be divided into slices, for example as illustrated in FIG. 3A. As illustrated in FIG. 3A, each slice  
15 may be divided into multiple tiles (four, in this example), each tile including multiple blocks. FIG. 3B illustrates an example tile that includes four 128x128 blocks. However, blocks of other sizes (e.g. 64x64, 32x32, etc.) may be used in some embodiments, and a tile may include more or fewer blocks.

**[0055]** Pre-filter 222 and wavelet transform 224 components of the video encoding system 220  
20 may then process each tile prior to encoding 226. In some embodiments, video encoding system 220 may include a separate pre-filter 222 component and wavelet transform 224 component for processing each tile. In this example, pre-filter 222A component and wavelet transform 224A component process tile 0, pre-filter 222B component and wavelet transform 224B component process tile 1, pre-filter 222C component and wavelet transform 224C component process tile 2,  
25 and pre-filter 222D component and wavelet transform 224D component process tile 3. The pre-filter 222 components perform pre-filtering of the tiles as described herein, and the wavelet transform 224 components decompose the tiles into frequency bands as described herein. However, in some embodiments, video encoding system 220 may include a single pre-filter 222 component and a single wavelet transform 224 component that process the tiles. In some  
30 embodiments, video encoding system 220 may include multiple (e.g., 2) pre-filter 222 components and multiple (e.g., 2) wavelet transform 224 components that each process multiple (e.g., 2) tiles.

**[0056]** Two encoders 226A and 226B may operate on two tiles from each slice (e.g., encoder 226A operates on tile 0 and tile 1; encoder 226B operates on tile 2 and tile 3). Each encoder 226 may multiplex the processing of blocks from different frequency bands (i.e., the 16 32x32 blocks  
35 illustrated in FIG. 3C) between its two tiles to allow for 16 time units between the processing of

blocks from the same frequency band. By multiplexing the processing of blocks between two tiles, dependencies between blocks in the same frequency band may be handled appropriately.

[0057] While embodiments are described in which each slice is divided into four tiles and two encoders operate on respective tiles from each slice, slices may be divided into more tiles (e.g., six or eight tiles) in some embodiments, and more encoders (e.g., three, four, or more encoders) may be used in some embodiments.

[0058] FIG. 3C illustrates performing a wavelet transform of a pixel block that stores pixel data to generate frequency band data prior to encoding, according to at least some embodiments. In this example, a two-level wavelet decomposition is applied by the wavelet transform 324 component to a 128x128 pixel block 300 to generate sixteen 32x32 blocks 302 of frequency data representing seven frequency bands. The frequency blocks 302 are then provided to an encoder 326 for encoding. For example, the frequency blocks 302 may be written to a buffer by the wavelet transform 324 component, and read from the buffer by the encoder 326 component.

[0059] In the labels of the frequency blocks 302, the letter L represents a low pass filter, and the letter H represents a high pass filter. The blocks 302 labeled with two letters represent a one-level (2D) wavelet transform or decomposition. In the blocks 302 labeled with two letters (representing three of the seven frequency bands LH, HL, and HH), the first letter represents a vertical filter (either high or low) performed first, and the second letter represents a horizontal filter (either high or low) performed second. The blocks 302 labeled with four letters represent a two-level wavelet transform or decomposition. In the blocks 302 labeled with four letters, the first two letters (LL) indicate that there was first a vertical low pass filter followed by a horizontal low pass filter; the second two letters indicate that the resulting LL block was then filtered four ways, LL, LH, HL, and HH (thus generating four of the seven frequency bands (LLLL, LLLH, LLHL, and LLHH)).

[0060] Decomposing the pixel data into frequency bands as illustrated in FIG. 3C allows the frequency bands to be buffered and processed as separate streams by the encoder 326. Processing the frequency bands as separate streams allows the encoder 326 component to multiplex the processing of the independent streams. In block-based encoding methods such as HEVC encoding, blocks (referred to as coding tree units (CTUs)) are processed in a block processing pipeline at multiple stages; two or more blocks may be at different stages of the pipeline at a given clock cycle, and the blocks move through the pipeline as the clock cycles. The processing of a given block may have dependencies on one or more previously processed neighbor blocks, for example one or more blocks in the row above the given block and/or the block to the left of the given block. By multiplexing the processing of the streams of frequency band data, the encoder 326 spaces out the processing of the blocks in a given stream, thus providing additional clock cycles to process a

neighbor block on which a given block has dependencies. For example, the block to the left of the given block may be several stages ahead of the given block in the pipeline when the given block reaches a stage that depends on the previously processed neighbor block. This allows the encoder 326 to better handle dependencies on previously processed blocks, and reduces or eliminates the need to wait for completion of processing of a neighbor block in the pipeline before processing the given block at a stage that depends on the neighbor block.

**[0061]** In addition, decomposing the pixel data into frequency bands as illustrated in FIG. 3C allows the frequency bands to be prioritized by the encoder 326 and the wireless interface. Typically, in image and video transmission, the lower frequencies are more important, while the higher frequencies are less important. Higher frequencies usually correspond to details in the image, and thus can be considered as lower priority. The higher frequency bands contain a smaller percentage of the energy in the image. Most of the energy is contained in the lower frequency bands. Decomposing the pixel data into frequency bands thus provides a priority ordering to the data stream that can be leveraged by the encoder 326 and the wireless interface when encoding and transmitting the data stream. For example, in some embodiments, different compression techniques may be used on the different frequency bands, with more aggressive compression applied to the lower priority bands, and more conservative compression applied to the higher priority bands. As another example, the priority ordering of the frequency bands may help in providing graceful degradation of the VR/MR system. Performance of the wireless connection can be monitored, and feedback from the device may be considered, to track performance of the overall system. If the system is falling behind for some reason, for example if the wireless connection degrades and bandwidth capacity of the wireless connection drops below a threshold, the encoder 326 and wireless interface may prioritize the encoding and transmission of one or more of the lower frequency bands, and may reduce or drop the encoding and/or transmission of one or more of the frequency levels that have been assigned a lower priority level, for example one or more of the higher frequency bands.

**[0062]** As described above, the wavelet transform decomposes an image into frequency bands. In some embodiments, this may be leveraged to send the same signal to displays of varying resolution. As an example, a two-level wavelet decomposition may be applied to decompose the signal into seven frequency bands. If four of the bands (LLLL, LLLH, LLHL and LLHH) are sent to a display panel, the bands may be reconstructed to the original intended resolution at lower visual quality. As an alternative, the bands may be reconstructed at 1/4th resolution (1/2 in each dimension) which may be suitable for a display panel with lower display resolution.

**[0063]** FIG. 4 is a high-level flowchart of a method of operation for VR/MR systems that include video encoding systems as illustrated in FIGS. 1 and 2, according to at least some

embodiments. As indicated at 400, the device sends data to the base station over the wireless connection. As indicated at 410, the base station renders frames including virtual content based at least in part on the data received from the device. As indicated at 420, the base station compresses the rendered data and sends the compressed data to the device over the wireless connection. As indicated at 430, the device decompresses and displays the virtual content to generate a 3D virtual view for viewing by the user. As indicated by the arrow returning from 430 to 400, the method continues as long as the user is using the VR/MR system.

**[0064]** In some embodiments, rather than rendering and encoding entire frames in the base station and transmitting the rendered frames to the device, the base station may render and encode parts of frames (referred to as slices) and transmit the encoded slices to the device as they are ready. A slice may, for example, be a row of 128x128 blocks. In some embodiments, the video encoding system may perform tile-based rendering, encoding, and transmittal. In tile-based rendering, encoding, and transmittal, each slice may be divided into multiple tiles each including one or more blocks (e.g., four tiles, each including four blocks), and the base station may render and encode the tiles and transmit the encoded tiles to the device as they are ready.

**[0065]** FIG. 5 is a flowchart of a method of operation for a video encoding system as illustrated in FIG. 1, according to at least some embodiments. The method of FIG. 5 may, for example, be performed at 420 of FIG. 4. The method of FIG. 5 assumes slice-based encoding and transmission is being performed. However, in some embodiments, tile-based encoding and transmission may be performed.

**[0066]** As indicated at 510, the pre-filter component applies lens warp and/or foveation filters to pixel blocks in a slice of the frame. In some embodiments, pre-filtering may include performing a lens warp on the frames on the base station prior to the wavelet transform. The lens warp is performed to correct for the distortion of the images introduced by the lenses on the device that the images are viewed through, thus improving quality of the images. In some embodiments, the device may store lens warp data for the lenses, for example generated by a calibration process, and may provide the lens warp data to the base station over the wireless connection. The base station may then perform the lens warp on the frames based on the lens warp data for that device. Performing the lens warp on the base station in the pre-filter stage may reduce the resolution of the frames prior to performing the wavelet transform and encoding, which may help in improving compression, and may reduce latency and bandwidth usage on the wireless link. In addition, by performing the lens warp on the base station in the pre-filter stage rather than on the device after decoding, filtering of the image data may only need to be performed once, as opposed to performing filtering on the base station to reduce resolution prior to encoding and then performing lens warp filtering on the device.

**[0067]** In some embodiments, pre-filtering at 510 may also include filtering to reduce resolution in peripheral regions while maintaining higher resolution in foveated regions. In some embodiments, gaze tracking information obtained from the device may be used to identify the direction in which the user is currently looking. A region of the frame that corresponds to the fovea (referred to as the foveated region) may be identified based at least in part on the determined gaze direction. The peripheral region (i.e. the portion of the frame outside the foveated region) may be pre-filtered to reduce information based on knowledge of the human vision system, for example by filtering high frequency information and/or increasing color compression. Pre-filtering of the peripheral region may help to provide improved compression of the frame.

**[0068]** As indicated at 520, the wavelet transform component applies a wavelet transform technique to the pixel blocks to decompose the pixel data into N (e.g., 7) frequency bands. The frequency bands are then organized into blocks that are provided to a block-based encoder for encoding/compression. As an example, a frame may be divided into 128x128 blocks, and a two-level wavelet decomposition may be applied to each 128x128 block to generate 16 32x32 blocks of frequency data representing seven frequency bands, for example as illustrated in FIG. 3C.

**[0069]** As indicated at 530, the encoder applies an encoding technique to the frequency bands in the blocks to compress the data. The encoder may, for example, be a High Efficiency Video Coding (HEVC) encoder. However, other encoding techniques may be used in some embodiments. Decomposing the pixel data into frequency bands as indicated at element 520 allows the frequency bands to be buffered and processed as separate streams by the encoder. Processing the frequency bands as separate streams allows the encoder component to multiplex the processing of the independent streams. In block-based encoding methods such as HEVC encoding, blocks (referred to as coding tree units (CTUs)) are processed in a pipeline at multiple stages; two or more blocks may be at different stages of the pipeline at a given clock cycle, and the blocks move through the pipeline as the clock cycles. The processing of a given block may have dependencies on one or more previously processed neighbor blocks, for example one or more blocks in the row above the given block and/or the block to the left of the given block. By multiplexing the processing of the streams, the encoder spaces out the processing of the blocks in a given stream, thus providing additional clock cycles to process a neighbor block on which a given block has dependencies. For example, the block to the left of the given block may be several stages ahead of the given block in the pipeline when the given block reaches a stage that depends on the previously processed neighbor block. This allows the encoder to better handle dependencies on previously processed blocks, and reduces or eliminates the need to wait for completion of processing of a neighbor block in the pipeline before processing the given block at a stage that depends on the neighbor block.

[0070] As indicated at 540, the wireless interface packetizes the compressed data and sends the packets to the device over the wireless connection.

[0071] Decomposing the pixel data into frequency bands as indicated at element 520 allows the frequency bands to be prioritized by the encoder at element 530 and the wireless interface at element 540. Typically, in image and video transmission, the lower frequencies are more important, while the higher frequencies are less important. Higher frequencies usually correspond to details in the image, and thus can be considered as lower priority. The higher frequency bands contain a smaller percentage of the energy in the image. Most of the energy is contained in the lower frequency bands. Decomposing the pixel data into frequency bands thus provides a priority ordering to the data stream that can be leveraged by the encoder and the wireless interface when encoding and transmitting the data stream. For example, in some embodiments, different compression techniques may be used on the different frequency bands, with more aggressive compression applied to the lower priority bands, and more conservative compression applied to the higher priority bands. As another example, the priority ordering of the frequency bands may help in providing graceful degradation of the VR/MR system. Performance of the wireless connection can be monitored, and feedback from the device may be considered, to track performance of the overall system. If the system is falling behind for some reason, for example if the wireless connection degrades and bandwidth capacity of the wireless connection drops below a threshold, the encoder and wireless interface may prioritize the encoding and transmission of one or more of the lower frequency bands, and may reduce or drop the encoding and/or transmission of one or more of the frequency levels that have been assigned a lower priority level, for example one or more of the higher frequency bands.

[0072] At 550, if there are more slices to be encoded and transmitted, the method returns to element 510 to process the next slice. Otherwise, at 560, if there are more frames to be encoded and transmitted, the method returns to element 510 to begin processing the next frame.

[0073] FIG. 6 is a flowchart of a method of operation for a video encoding system as illustrated in FIG. 2, according to at least some embodiments. The method of FIG. 6 may, for example, be performed at 420 of FIG. 4. In the method of FIG. 6, the video encoding system may perform tile-based encoding using multiple encoders to process respective tiles from each slice.

[0074] As indicated at 600, a rendering engine renders a slice including multiple tiles (four tiles, in this example), each tiles including multiple pixel blocks (four 128x128 pixel blocks, in this example).

[0075] As indicated at 610, the pre-filter component applies lens warp and/or foveation filters to the slice. In some embodiments, pre-filtering may include performing a lens warp on the frames on the base station prior to the wavelet transform. The lens warp is performed to correct for the

distortion of the images introduced by the lenses on the device that the images are viewed through, thus improving quality of the images. In some embodiments, the device may store lens warp data for the lenses, for example generated by a calibration process, and may provide the lens warp data to the base station over the wireless connection. The base station may then perform the lens warp on the frames based on the lens warp data for that device. Performing the lens warp on the base station in the pre-filter stage may reduce the resolution of the frames prior to performing the wavelet transform and encoding, which may help in improving compression, and may reduce latency and bandwidth usage on the wireless link. In addition, by performing the lens warp on the base station in the pre-filter stage rather than on the device after decoding, filtering of the image data may only need to be performed once, as opposed to performing filtering on the base station to reduce resolution prior to encoding and then performing lens warp filtering on the device.

**[0076]** In some embodiments, pre-filtering at 610 may also include filtering to reduce resolution in peripheral regions while maintaining higher resolution in foveated regions. In some embodiments, gaze tracking information obtained by the device may be used to identify the direction in which the user is currently looking. A region of the frame that corresponds to the fovea (referred to as the foveated region) may be identified based at least in part on the determined gaze direction. The peripheral region (i.e. the portion of the frame outside the foveated region) may be pre-filtered to reduce information based on knowledge of the human vision system, for example by filtering high frequency information and/or increasing color compression. Pre-filtering of the peripheral region may help to provide improved compression of the frame.

**[0077]** In some embodiments, the video encoding system may include a single pre-filter component that process the tiles. In some embodiments, the video encoding system may include a separate pre-filter component for processing each tile. In some embodiments, the video encoding system may include multiple (e.g., 2) pre-filter components that each process multiple (e.g., 2) tiles.

**[0078]** As indicated at 620, the wavelet transform component applies a wavelet transform technique to the pixel blocks in the slice to decompose the pixel data into N (e.g., 7) frequency bands. The frequency bands are then organized into blocks (e.g., CTUs) that can be provided to a block-based encoder for encoding/compression. As an example, a frame may be divided into 128x128 blocks, and a two-level wavelet decomposition may be applied to each 128x128 block to generate 16 32x32 blocks of frequency data representing seven frequency bands, for example as illustrated in FIG. 3C.

**[0079]** In some embodiments, the video encoding system may include a single wavelet transform component that process the tiles. In some embodiments, the video encoding system may include a separate wavelet transform component for processing each tile. In some embodiments,

the video encoding system may include multiple (e.g., 2) wavelet transform components that each process multiple (e.g., 2) tiles.

**[0080]** In the method of FIG. 6, an example video encoding system includes two encoders configured to encode the blocks of frequency data from different ones of the slices that are generated at element 620 in parallel. For example, a first encoder may be configured to encode blocks from slices 0 and 1, and a second encoder may be configured to encode blocks from slices 2 and 3. As indicated at 630A, the first encoder applies an encoding technique to the frequency bands in the blocks of tiles 0 and 1, multiplexing the processing of blocks from different frequency bands between the two different tiles. As indicated at 630B, a second encoder applies an encoding technique to the frequency bands in the blocks of tiles 2 and 3, multiplexing the processing of blocks from different frequency bands between the two different tiles.

**[0081]** The encoders may, for example, be a High Efficiency Video Coding (HEVC) encoders. However, other encoding techniques may be used in some embodiments. Decomposing the pixel data into frequency bands as indicated at element 620 allows the frequency bands to be buffered and processed as separate streams by the encoders at elements 630A and 630B. Processing the frequency bands as separate streams allows the encoders to multiplex the processing of the independent streams. In block-based encoding methods such as HEVC encoding, blocks (referred to as coding tree units (CTUs)) are processed in a pipeline at multiple stages; two or more blocks may be at different stages of the pipeline at a given clock cycle, and the blocks move through the pipeline as the clock cycles. The processing of a given block may have dependencies on one or more previously processed neighbor blocks, for example one or more blocks in the row above the given block and/or the block to the left of the given block. By multiplexing the processing of the streams, the encoder spaces out the processing of the blocks in a given stream, thus providing additional clock cycles to process a neighbor block on which a given block has dependencies. For example, the block to the left of the given block may be several stages ahead of the given block in the pipeline when the given block reaches a stage that depends on the previously processed neighbor block. This allows the encoder to better handle dependencies on previously processed blocks, and reduces or eliminates the need to wait for completion of processing of a neighbor block in the pipeline before processing the given block at a stage that depends on the neighbor block.

**[0082]** As indicated at 640, the wireless interface packetizes the compressed data generated by the encoders at element 530 and sends the packets to the device over the wireless connection. As indicated by the arrow returning from 650 to 600, the method continues as long as the user is using the VR/MR system.

**[0083]** Decomposing the pixel data into frequency bands as indicated at element 620 allows the frequency bands to be prioritized by the encoders at elements 630A and 630B and the wireless

interface at element 640. Typically, in image and video transmission, the lower frequencies are more important, while the higher frequencies are less important. Higher frequencies usually correspond to details in the image, and thus can be considered as lower priority. The higher frequency bands contain a smaller percentage of the energy in the image. Most of the energy is contained in the lower frequency bands. Decomposing the pixel data into frequency bands thus provides a priority ordering to the data stream that can be leveraged by the encoder and the wireless interface when encoding and transmitting the data stream. For example, in some embodiments, different compression techniques may be used on the different frequency bands, with more aggressive compression applied to the lower priority bands, and more conservative compression applied to the higher priority bands. As another example, the priority ordering of the frequency bands may help in providing graceful degradation of the VR/MR system. Performance of the wireless connection can be monitored, and feedback from the device may be considered, to track performance of the overall system. If the system is falling behind for some reason, for example if the wireless connection degrades and bandwidth capacity of the wireless connection drops below a threshold, the encoder and wireless interface may prioritize the encoding and transmission of one or more of the lower frequency bands, and may reduce or drop the encoding and/or transmission of one or more of the frequency levels that have been assigned a lower priority level, for example one or more of the higher frequency bands.

**[0084]** While embodiments are described in which each slice is divided into four tiles and two encoders operate on respective tiles from each slice, slices may be divided into more tiles (e.g., six or eight tiles) in some embodiments, and more encoders (e.g., three or four encoders) may be used in some embodiments.

#### Example VR/MR system

**[0085]** FIG. 7 illustrates an example VR/MR system 2000 that may implement a video encoding system, according to at least some embodiments. A VR/MR system 2000 may include at least one device 2150 (e.g., a notebook or laptop computer, pad or tablet device, smartphone, hand-held computing device or an HMD such as a headset, helmet, goggles, or glasses that may be worn by a user) and a computing device 2100 (referred to herein as a base station). The base station 2100 renders VR or MR frames including virtual content, encodes the frames, and transmits the encoded frames over a wireless connection 2180 to the device 2150 for decoding and display by the device 2150.

**[0086]** The base station 2100 and device 2150 may each include wireless communications technology that allows the base station 2100 and device 2150 to communicate and exchange data

via the wireless connection 2180. In some embodiments, the wireless connection 2180 may be implemented according to a proprietary wireless communications technology that provides a highly directional wireless link between the device 2150 and the base station 2100. However, other commercial (e.g., Wi-Fi, Bluetooth, etc.) or proprietary wireless communications technologies may be used in some embodiments.

**[0087]** In some embodiments, the device 2150 may include sensors that collect information about the user's environment (e.g., video, depth information, lighting information, etc.) and/or about the user (e.g., the user's expressions, eye movement, gaze direction, hand gestures, etc.). The device 2150 may transmit at least some of the information collected by sensors to the base station 2100 via wireless connection 2180. The base station 2100 may render frames for display by the device 2150 that include virtual content based at least in part on the various information obtained from the sensors, encode the frames, and transmit the encoded frames to the device 2150 for decoding and display to the user via the wireless connection 2180. To encode and transmit the frames, the base station 2100 may implement a video encoding system as illustrated in FIGS. 1 through 6.

**[0088]** FIG. 8 is a block diagram illustrating functional components of and processing in an example VR/MR system as illustrated in FIG. 7, according to some embodiments. Device 2150 may be, but is not limited to, a notebook or laptop computer, pad or tablet device, smartphone, hand-held computing device or an HMD such as a headset, helmet, goggles, or glasses that may be worn by a user. Device 2150 may include a display 2156 component or subsystem that may implement any of various types of virtual or augmented reality display technologies. For example, an HMD device 2150 may be a near-eye system that displays left and right images on screens in front of the user's eyes that are viewed by a subject, such as DLP (digital light processing), LCD (liquid crystal display) and LCoS (liquid crystal on silicon) technology VR systems. As another example, an HMD device 2150 may be a direct retinal projector system that scans left and right images, pixel by pixel, to the subject's eyes. To scan the images, left and right projectors generate beams that are directed to left and right reflective components (e.g., ellipsoid mirrors) located in front of the user's eyes; the reflective components reflect the beams to the user's eyes. To create a three-dimensional (3D) effect, virtual content at different depths or distances in the 3D virtual view are shifted left or right in the two images as a function of the triangulation of distance, with nearer objects shifted more than more distant objects.

**[0089]** Device 2150 may also include a controller 2154 configured to implement device -side functionality of the VR/MR system 2000 as described herein. In some embodiments, HMD 2150 may also include memory 2170 configured to store software (code 2172) of the device component of the VR/MR system 2000 that is executable by the controller 2154, as well as data 2174 that may

be used by the software when executing on the controller 2154. In various embodiments, the controller 2154 may be a uniprocessor system including one processor, or a multiprocessor system including several processors (e.g., two, four, eight, or another suitable number). The controller 2154 may include central processing units (CPUs) configured to implement any suitable instruction set architecture, and may be configured to execute instructions defined in that instruction set architecture. For example, in various embodiments the controller 2154 may include general-purpose or embedded processors implementing any of a variety of instruction set architectures (ISAs), such as the x86, PowerPC, SPARC, RISC, or MIPS ISAs, or any other suitable ISA. In multiprocessor systems, each of the processors may commonly, but not necessarily, implement the same ISA. The controller 2154 may employ any microarchitecture, including scalar, superscalar, pipelined, superpipelined, out of order, in order, speculative, non-speculative, etc., or combinations thereof. The controller 2154 may include circuitry to implement microcoding techniques. The controller 2154 may include one or more processing cores each configured to execute instructions. The controller 2154 may include one or more levels of caches, which may employ any size and any configuration (set associative, direct mapped, etc.). In some embodiments, the controller 2154 may include at least one graphics processing unit (GPU), which may include any suitable graphics processing circuitry. Generally, a GPU may be configured to render objects to be displayed into a frame buffer (e.g., one that includes pixel data for an entire frame). A GPU may include one or more graphics processors that may execute graphics software to perform a part or all of the graphics operation, or hardware acceleration of certain graphics operations. In some embodiments, the controller 2154 may include one or more other components for processing and rendering video and/or images, for example image signal processors (ISPs), encoder/decoders (codecs), etc. In some embodiments, controller 2154 may include at least one system on a chip (SOC).

**[0090]** The memory 2170 may include any type of memory, such as dynamic random access memory (DRAM), synchronous DRAM (SDRAM), double data rate (DDR, DDR2, DDR3, etc.) SDRAM (including mobile versions of the SDRAMs such as mDDR3, etc., or low power versions of the SDRAMs such as LPDDR2, etc.), RAMBUS DRAM (RDRAM), static RAM (SRAM), etc. In some embodiments, one or more memory devices may be coupled onto a circuit board to form memory modules such as single inline memory modules (SIMMs), dual inline memory modules (DIMMs), etc. Alternatively, the devices may be mounted with an integrated circuit implementing system in a chip-on-chip configuration, a package-on-package configuration, or a multi-chip module configuration.

**[0091]** In some embodiments, sensors 2160 may include, but are not limited to, one or more gaze tracking sensors (e.g., IR cameras with an IR illumination source) that may be used to track

position and movement of the user's eyes. In some embodiments, there may be two gaze tracking sensors, with each gaze tracking sensor tracking a respective eye. In some embodiments, the information collected by the gaze tracking sensors may be used to adjust the rendering of images by the base station 2100, and/or to adjust the projection of the images by the projection system of the device 2150, based on the direction and angle at which the user's eyes are looking. For example, in some embodiments, content of the images in a region around the location at which the user's eyes are currently looking may be rendered with more detail and at a higher resolution than content in regions at which the user is not looking, which allows available processing time for image data to be spent on content viewed by the foveal regions of the eyes rather than on content viewed by the peripheral regions of the eyes. Similarly, content of images in regions at which the user is not looking may be compressed more than content of the region around the point at which the user is currently looking. In some embodiments there may be two gaze tracking sensors located on an inner surface of the device 2150 at positions such that the sensors have views of respective ones of the user's eyes. However, in various embodiments, more or fewer gaze tracking sensors may be used, and gaze tracking sensors may be positioned at other locations. In an example non-limiting embodiment, each gaze tracking sensor may include an IR light source and IR camera, for example a 400x400 pixel count camera with a frame rate of 120 FPS or greater, HFOV of 70 degrees, and with a working distance of 10 millimeters (mm) to 80mm.

**[0092]** In some embodiments, the device 2150 may include at least one inertial-measurement unit (IMU) 2162 configured to detect position, orientation, and/or motion of the device 2150, and to provide the detected position, orientation, and/or motion data to the controller 2154 of the HMD 2150 and/or to the base station 2100.

**[0093]** Device 2150 may also include a wireless interface 2152 configured to communicate with an external base station 2100 via a wireless connection 2180 to send sensor inputs to the base station 2100 and to receive compressed rendered frames, slices, or tiles from the base station 2100. In some embodiments, the wireless interface 2152 may implement a proprietary wireless communications technology that provides a highly directional wireless link between the device 2150 and the base station 2100. However, other commercial (e.g., Wi-Fi, Bluetooth, etc.) or proprietary wireless communications technologies may be used in some embodiments.

**[0094]** The base station 2100 may be an external device (e.g., a computing system, game console, etc.) that is communicatively coupled to device 2150 via a wireless interface 2180. The base station 2100 may include one or more of various types of processors (e.g., SOCs, CPUs, ISPs, GPUs, codecs, and/or other components) for rendering, filtering, encoding, and transmitting video and/or images. The base station 2100 may render frames (each frame including a left and right image) that include virtual content based at least in part on the various inputs obtained from the

sensors 2160 via the wireless connection 2180, filter and compress the rendered frames (or slices of the frames) using a video encoding system as described herein, and transmit the compressed frames or slices to the device 2150 for display.

**[0095]** Base station 2100 may be or may include any type of computing system or computing device, such as a desktop computer, notebook or laptop computer, pad or tablet device, smartphone, hand-held computing device, game controller, game system, and so on. Base station 2100 may include a controller 2110 comprising one or more processors that implement base-side functionality of the VR/MR system 2000 including the video encoding system as described herein. Base station 2100 may also include memory 2120 configured to store software (code 2122) of the base station component of the VR/MR system 2000 that is executable by the base station controller 2110, as well as data 2124 that may be used by the software when executing on the controller 2110.

**[0096]** In various embodiments, the base station controller 2110 may be a uniprocessor system including one processor, or a multiprocessor system including several processors (e.g., two, four, eight, or another suitable number). The controller 2110 may include central processing units (CPUs) configured to implement any suitable instruction set architecture and may be configured to execute instructions defined in that instruction set architecture. For example, in various embodiments the controller 2110 may include general-purpose or embedded processors implementing any of a variety of instruction set architectures (ISAs), such as the x86, PowerPC, SPARC, RISC, or MIPS ISAs, or any other suitable ISA. In multiprocessor systems, each of the processors may commonly, but not necessarily, implement the same ISA. The controller 2110 may employ any microarchitecture, including scalar, superscalar, pipelined, superpipelined, out of order, in order, speculative, non-speculative, etc., or combinations thereof. Controller 2110 may include circuitry to implement microcoding techniques. The controller 2110 may include one or more processing cores each configured to execute instructions. The controller 2110 may include one or more levels of caches, which may employ any size and any configuration (set associative, direct mapped, etc.). In some embodiments, the controller 2110 may include at least one graphics processing unit (GPU), which may include any suitable graphics processing circuitry. Generally, a GPU may be configured to render objects to be displayed into a frame buffer (e.g., one that includes pixel data for an entire frame). A GPU may include one or more graphics processors that may execute graphics software to perform a part or all of the graphics operation, or hardware acceleration of certain graphics operations. In some embodiments, the controller 2110 may include one or more other components for processing, rendering, filtering, and encoding video and/or images as described herein, for example one or more of various types of integrated circuits (ICs),

image signal processors (ISPs), encoder/decoders (codecs), etc. In some embodiments, the controller 2110 may include at least one system on a chip (SOC).

[0097] The base station memory 2120 may include any type of memory, such as dynamic random access memory (DRAM), synchronous DRAM (SDRAM), double data rate (DDR, DDR2, DDR3, etc.) SDRAM (including mobile versions of the SDRAMs such as mDDR3, etc., or low power versions of the SDRAMs such as LPDDR2, etc.), RAMBUS DRAM (RDRAM), static RAM (SRAM), etc. In some embodiments, one or more memory devices may be coupled onto a circuit board to form memory modules such as single inline memory modules (SIMMs), dual inline memory modules (DIMMs), etc. Alternatively, the devices may be mounted with an integrated circuit implementing system in a chip-on-chip configuration, a package-on-package configuration, or a multi-chip module configuration.

[0098] Base station 2100 may also include one or more wireless technology interfaces 2130 configured to communicate with device 2150 via a wireless connection 2180 to receive sensor inputs from the device 2150 and send compressed frames, slices, or tiles from the base station 2100 to the device 2150. In some embodiments, a wireless technology interface 2130 may implement a proprietary wireless communications technology that provides a highly directional wireless link between the device 2150 and the base station 2100. In some embodiments, the directionality and band width of the wireless communication technology may support multiple devices 2150 communicating with the base station 2100 at the same time to thus enable multiple users to use the system 2000 at the same time in a co-located environment. However, other commercial (e.g., Wi-Fi, Bluetooth, etc.) or proprietary wireless communications technologies may be used in some embodiments.

[0099] In some embodiments, the base station 2100 may be configured to render and transmit frames to the device 2150 to provide a 3D virtual view for the user based at least in part on sensor 2160 inputs received from the device 2150. In some embodiments, the virtual view may include renderings of the user's environment, including renderings of real objects in the user's environment, based on video captured by one or more scene cameras (e.g., RGB (visible light) video cameras) that capture high-quality, high-resolution video of the user's environment in real time for display. In some embodiments, the virtual view may also include virtual content (e.g., virtual objects, virtual tags for real objects, avatars of the user, etc.) rendered and composited with the projected 3D view of the user's real environment by the base station 2100.

[0100] While not shown in FIGS. 7 and 8, in some embodiments the VR/MR system 2000 may include one or more other components. For example, the system may include a cursor control device (e.g., mouse) for moving a virtual cursor in the 3D virtual view to interact with virtual content. While FIGS. 7 and 8 show a single device 2150, in some embodiments the VR/MR

system 2000 may support multiple devices 2150 communicating with the base station 2100 at the same time to thus enable multiple users to use the system at the same time in a co-located environment.

**[0101]** The methods described herein may be implemented in software, hardware, or a combination thereof, in different embodiments. In addition, the order of the blocks of the methods may be changed, and various elements may be added, reordered, combined, omitted, modified, etc. Various modifications and changes may be made as would be obvious to a person skilled in the art having the benefit of this disclosure. The various embodiments described herein are meant to be illustrative and not limiting. Many variations, modifications, additions, and improvements are possible. Accordingly, plural instances may be provided for components described herein as a single instance. Boundaries between various components, operations and data stores are somewhat arbitrary, and particular operations are illustrated in the context of specific illustrative configurations. Other allocations of functionality are envisioned and may fall within the scope of claims that follow. Finally, structures and functionality presented as discrete components in the example configurations may be implemented as a combined structure or component. These and other variations, modifications, additions, and improvements may fall within the scope of embodiments as defined in the claims that follow.

**CLAIMS****WHAT IS CLAIMED IS:**

- 5           1.       A video encoding system, comprising:
- a wavelet transform component comprising one or more processors configured to:
- apply a wavelet transform to pixel blocks from a rendered video frame to
- decompose the pixel blocks into a plurality of frequency bands; and
- 10           organize the frequency bands from the pixel blocks into a plurality of frequency
- blocks each including frequency data for one of the frequency bands; and
- at least one block-based encoder configured to:
- apply an encoding technique to the frequency blocks to compress the frequency
- 15           data in the frequency blocks; and
- provide the compressed frequency data to a wireless interface for transmittal over
- a wireless connection to a receiving device.
2.       The video encoding system as recited in claim 1, further comprising a pre-filter
- 20           component comprising one or more processors configured to pre-filter the pixel blocks from the
- video frame prior to the wavelet transform.
3.       The video encoding system as recited in claim 2, wherein, to pre-filter the pixel
- blocks, the pre-filter component applies lens warp filter to the video frame to correct for distortion
- 25           introduced by lenses of the receiving device.
4.       The video encoding system as recited in claim 2, wherein, to pre-filter the pixel
- blocks, the pre-filter component applies at least one filter to a peripheral region of the video frame
- to reduce resolution in the peripheral region while maintaining higher resolution in a foveated
- 30           region of the video frame.
5.       The video encoding system as recited in claim 1, wherein the rendered video frame
- is divided into a plurality of slices, each slice including at least one row of pixel blocks, and
- wherein the video encoding system is configured to iteratively perform:
- 35           obtain a slice of the rendered video frame;

apply the wavelet transform to the pixel blocks in the slice to decompose the pixel blocks into a plurality of frequency bands organized into a plurality of frequency blocks each including frequency data for one of the frequency bands;

5 apply the encoding technique to the frequency blocks from the slice to compress the frequency data in the frequency blocks; and

provide the compressed frequency data from the slice to the wireless interface for transmittal over the wireless connection to the at least one receiving device.

6. The video encoding system as recited in claim 1, wherein the rendered video frame  
10 is divided into a plurality of slices, each slice including two or more tiles, each tile including two or more pixel blocks, and wherein the video encoding system comprises at least two block-based encoders, each encoder configured to apply the encoding technique to the frequency blocks from a subset of the tiles from each slice to compress the frequency data in the frequency blocks from the respective subset of the tiles.

15 7. The video encoding system as recited in claim 6, wherein each encoder is configured to multiplex the processing of the frequency blocks for the frequency bands between the tiles in its respective subset to allow for multiple time units between the processing of frequency blocks of each frequency band.

20 8. The video encoding system as recited in claim 1, wherein the wavelet transform is a two-level wavelet decomposition applied to each pixel block from the video frame.

25 9. The video encoding system as recited in claim 1, wherein the wavelet transform is an N-level wavelet decomposition applied to each pixel block from the video frame, where N is adjustable.

30 10. The video encoding system as recited in claim 1, wherein applying the wavelet transform comprises applying a coring function to one or more wavelet coefficients of the wavelet decomposition.

11. The video encoding system as recited in claim 1, wherein the pixel blocks are 128x128 pixel blocks, and wherein the wavelet transform decomposes each pixel block into sixteen 32x32 blocks of frequency data representing seven different frequency bands.

35

12. A method, comprising:

one or more processors configured to implement:

5           applying a wavelet transform to pixel blocks from a rendered video frame to  
          decompose the pixel blocks into a plurality of frequency bands;  
          organizing the frequency bands from the pixel blocks into a plurality of frequency  
          blocks each including frequency data for one of the frequency bands; and  
          applying an encoding technique to the frequency blocks to compress the frequency  
          data in the frequency blocks; and  
10           transmitting the compressed frequency data to a receiving device over a wireless  
          connection.

13. The method as recited in claim 12, further comprising pre-filtering the pixel blocks  
15           from the video frame prior to applying the wavelet transform, wherein pre-filtering comprises one  
          or more of:

          applying a lens warp filter to the video frame to correct for distortion introduced by lenses  
          of the receiving device; or  
          applying at least one filter to a peripheral region of the video frame to reduce resolution in  
          the peripheral region while maintaining higher resolution in a foveated region of  
20           the video frame.

14. The method as recited in claim 12, wherein the video frame is divided into a  
          plurality of slices, each slice including at least one row of pixel blocks, and wherein the one or  
          more processors perform said applying a wavelet transform, said organizing, said applying an  
25           encoding technique, and said transmitting for each slice of the video frame.

15. The method as recited in claim 12, wherein the video frame is divided into a  
          plurality of slices, each slice including two or more tiles, each tile including two or more pixel  
          blocks, and wherein the one or more processors perform said applying a wavelet transform, said  
30           organizing, said applying an encoding technique, and said transmitting for each tile of the video  
          frame.

16. The method as recited in claim 15, wherein said applying an encoding technique is  
          performed by at least two block-based encoders, each encoder applying the encoding technique to  
35           the frequency blocks from a subset of the tiles from each slice to compress the frequency data in

the frequency blocks from the respective subset of the tiles, the method further comprising each encoder multiplexing the processing of the frequency blocks for the frequency bands between the tiles in its respective subset to allow for multiple time units between the processing of frequency blocks of each frequency band.

5

17. The method as recited in claim 12, wherein the wavelet transform is a two-level wavelet decomposition applied to each pixel block from the video frame.

18. A system, comprising:

10

a device comprising one or more processors and a display subsystem;

a base station comprising:

one or more processors configured to apply a wavelet transform to blocks of pixel data from a rendered video frame to decompose the pixel data into a plurality of frequency bands and organize the frequency bands into a plurality of frequency blocks; and

15

at least one block-based encoder configured to apply an encoding technique to the frequency blocks to compress the frequency bands; and

a wireless interface configured to transmit the compressed frequency bands over a wireless connection to the device;

20

wherein the one or more processors of the device are configured to:

decompress the compressed frequency bands received from the base station;

perform wavelet synthesis on the decompressed frequency bands to reconstruct the pixel data for the video frame; and

25

provide the reconstructed pixel data to the display subsystem for display.

19. The system as recited in claim 18, wherein the one or more processors of the base station are further configured to pre-filter the pixel data of the video frame prior to the wavelet transform, wherein, to pre-filter the pixel data, the one or more processors apply a lens warp filter to the video frame to correct for distortion introduced by lenses of the device.

30

20. The system as recited in claim 18, wherein the device is a head-mounted display (HMD).

35




FIG. 1

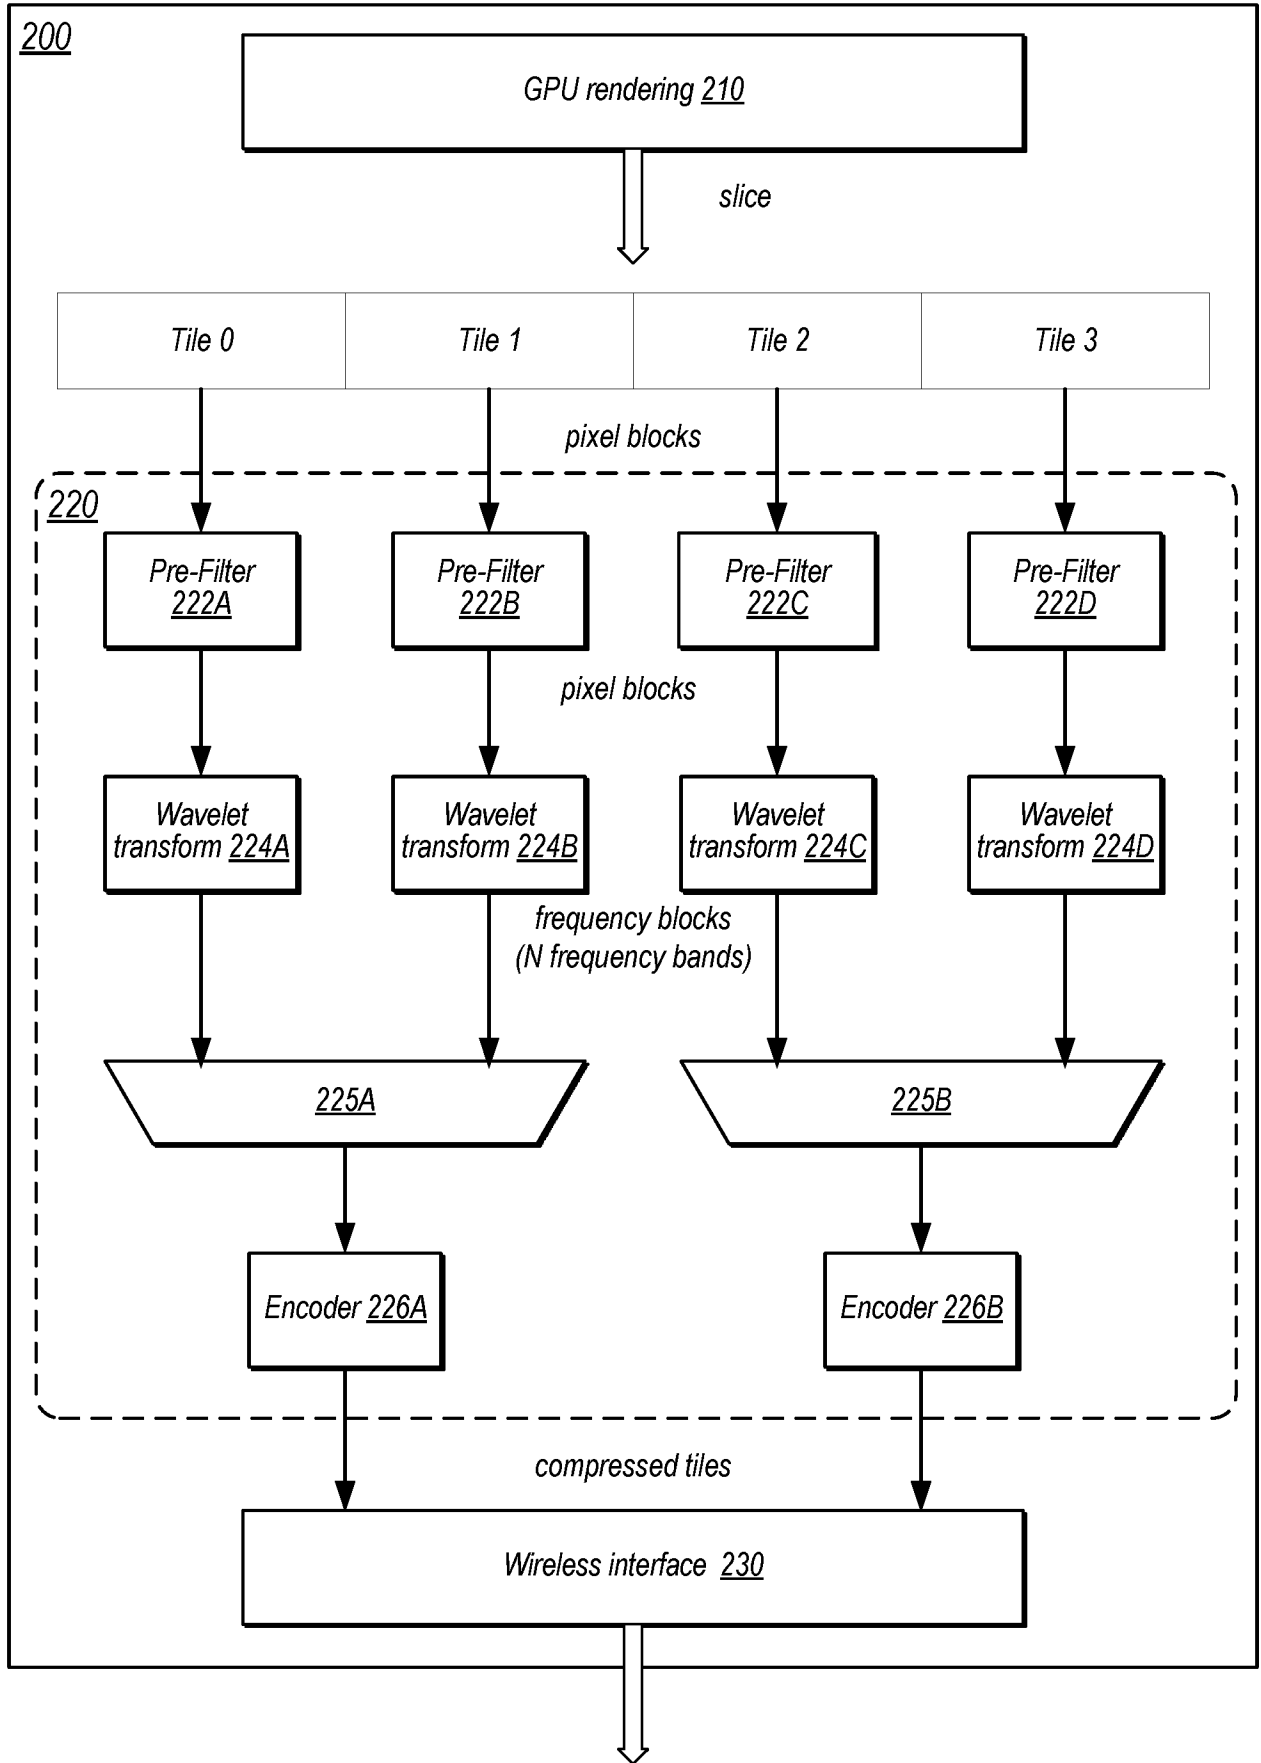


FIG. 2


3 / 9

Frame

<i>slice 0</i>	<i>Tile 0</i>	<i>Tile 1</i>	<i>Tile 2</i>	<i>Tile 3</i>
<i>slice 1</i>	<i>Tile 0</i>	<i>Tile 1</i>	<i>Tile 2</i>	<i>Tile 3</i>
<i>slice 2</i>	<i>Tile 0</i>	<i>Tile 1</i>	<i>Tile 2</i>	<i>Tile 3</i>
<i>slice 3</i>	<i>Tile 0</i>	<i>Tile 1</i>	<i>Tile 2</i>	<i>Tile 3</i>

⋮

**FIG. 3A**



**FIG. 3B**

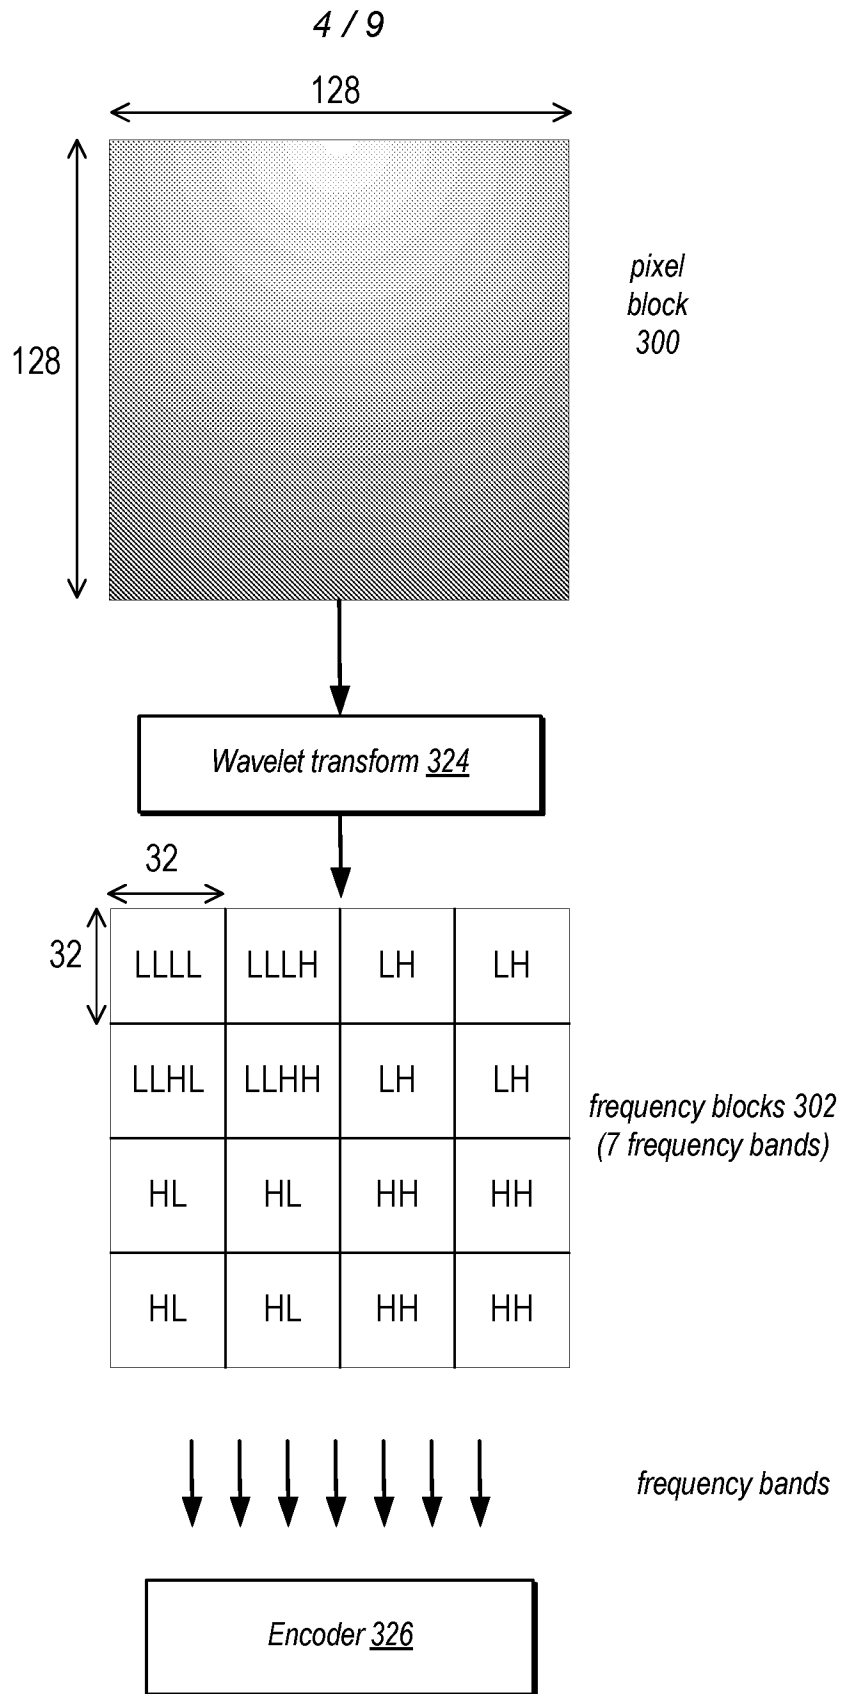


FIG. 3C

5 / 9

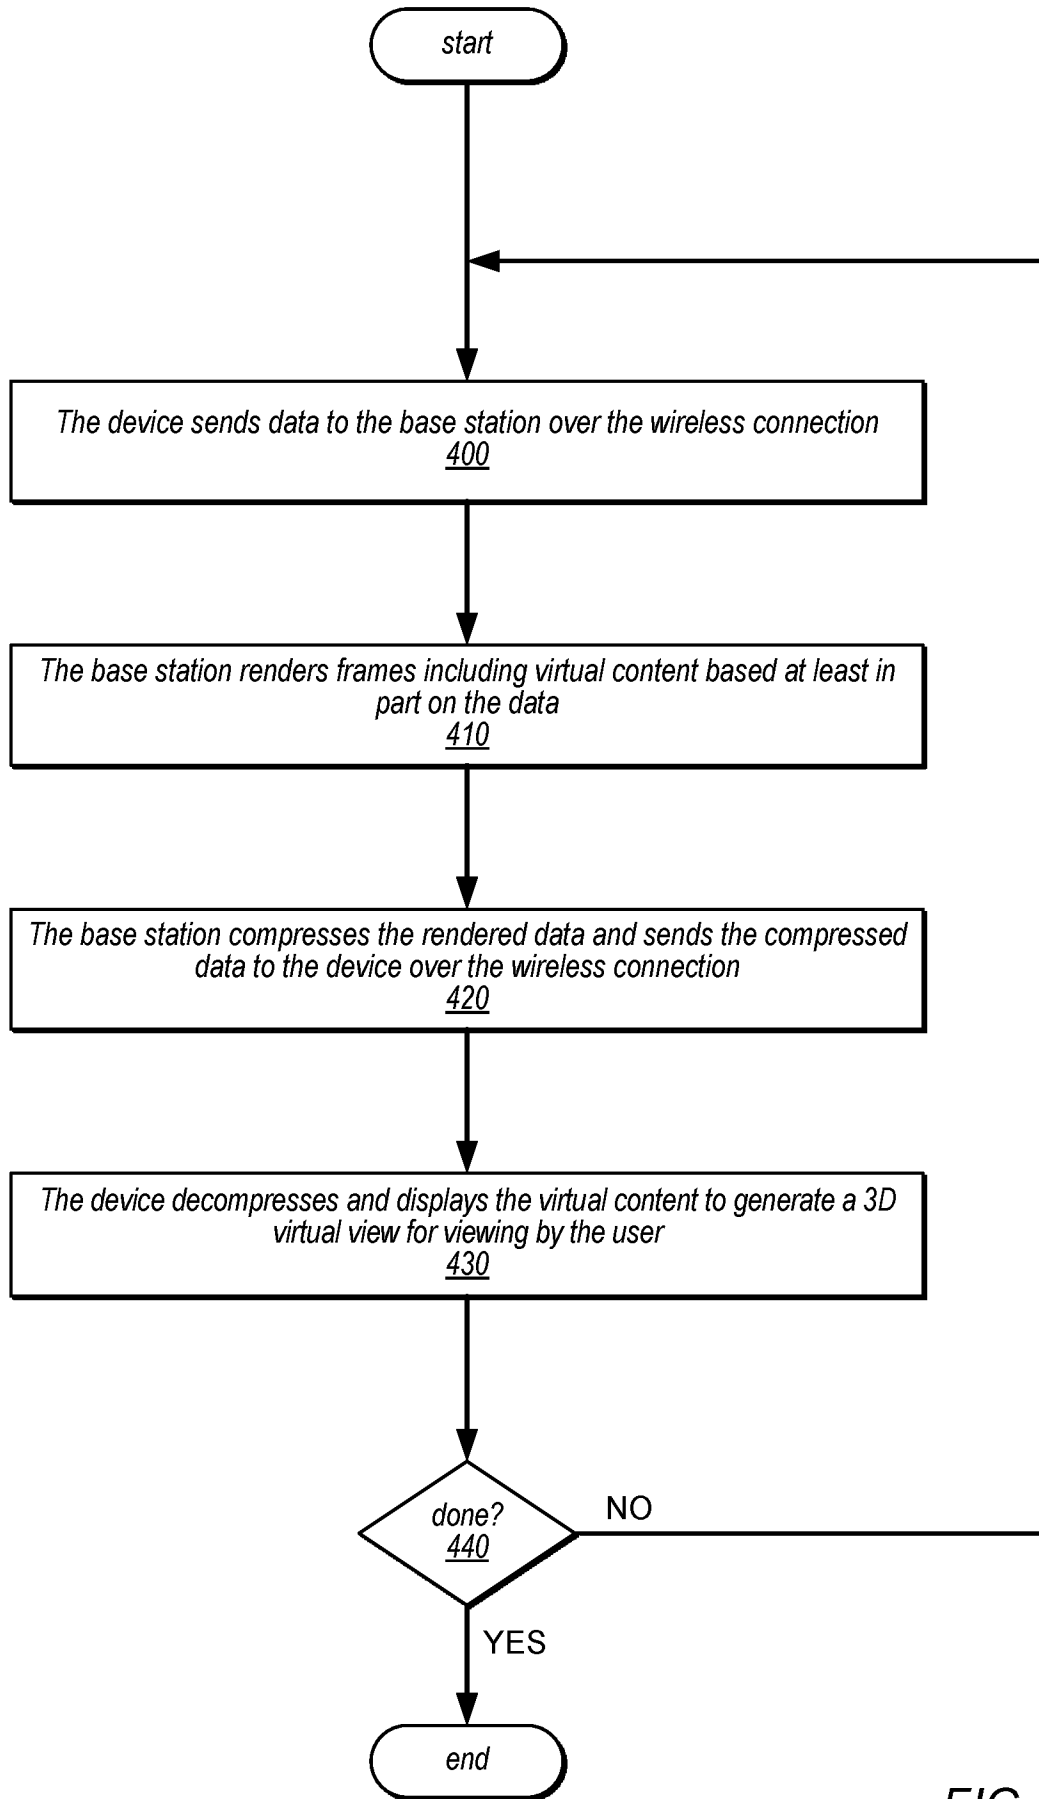


FIG. 4

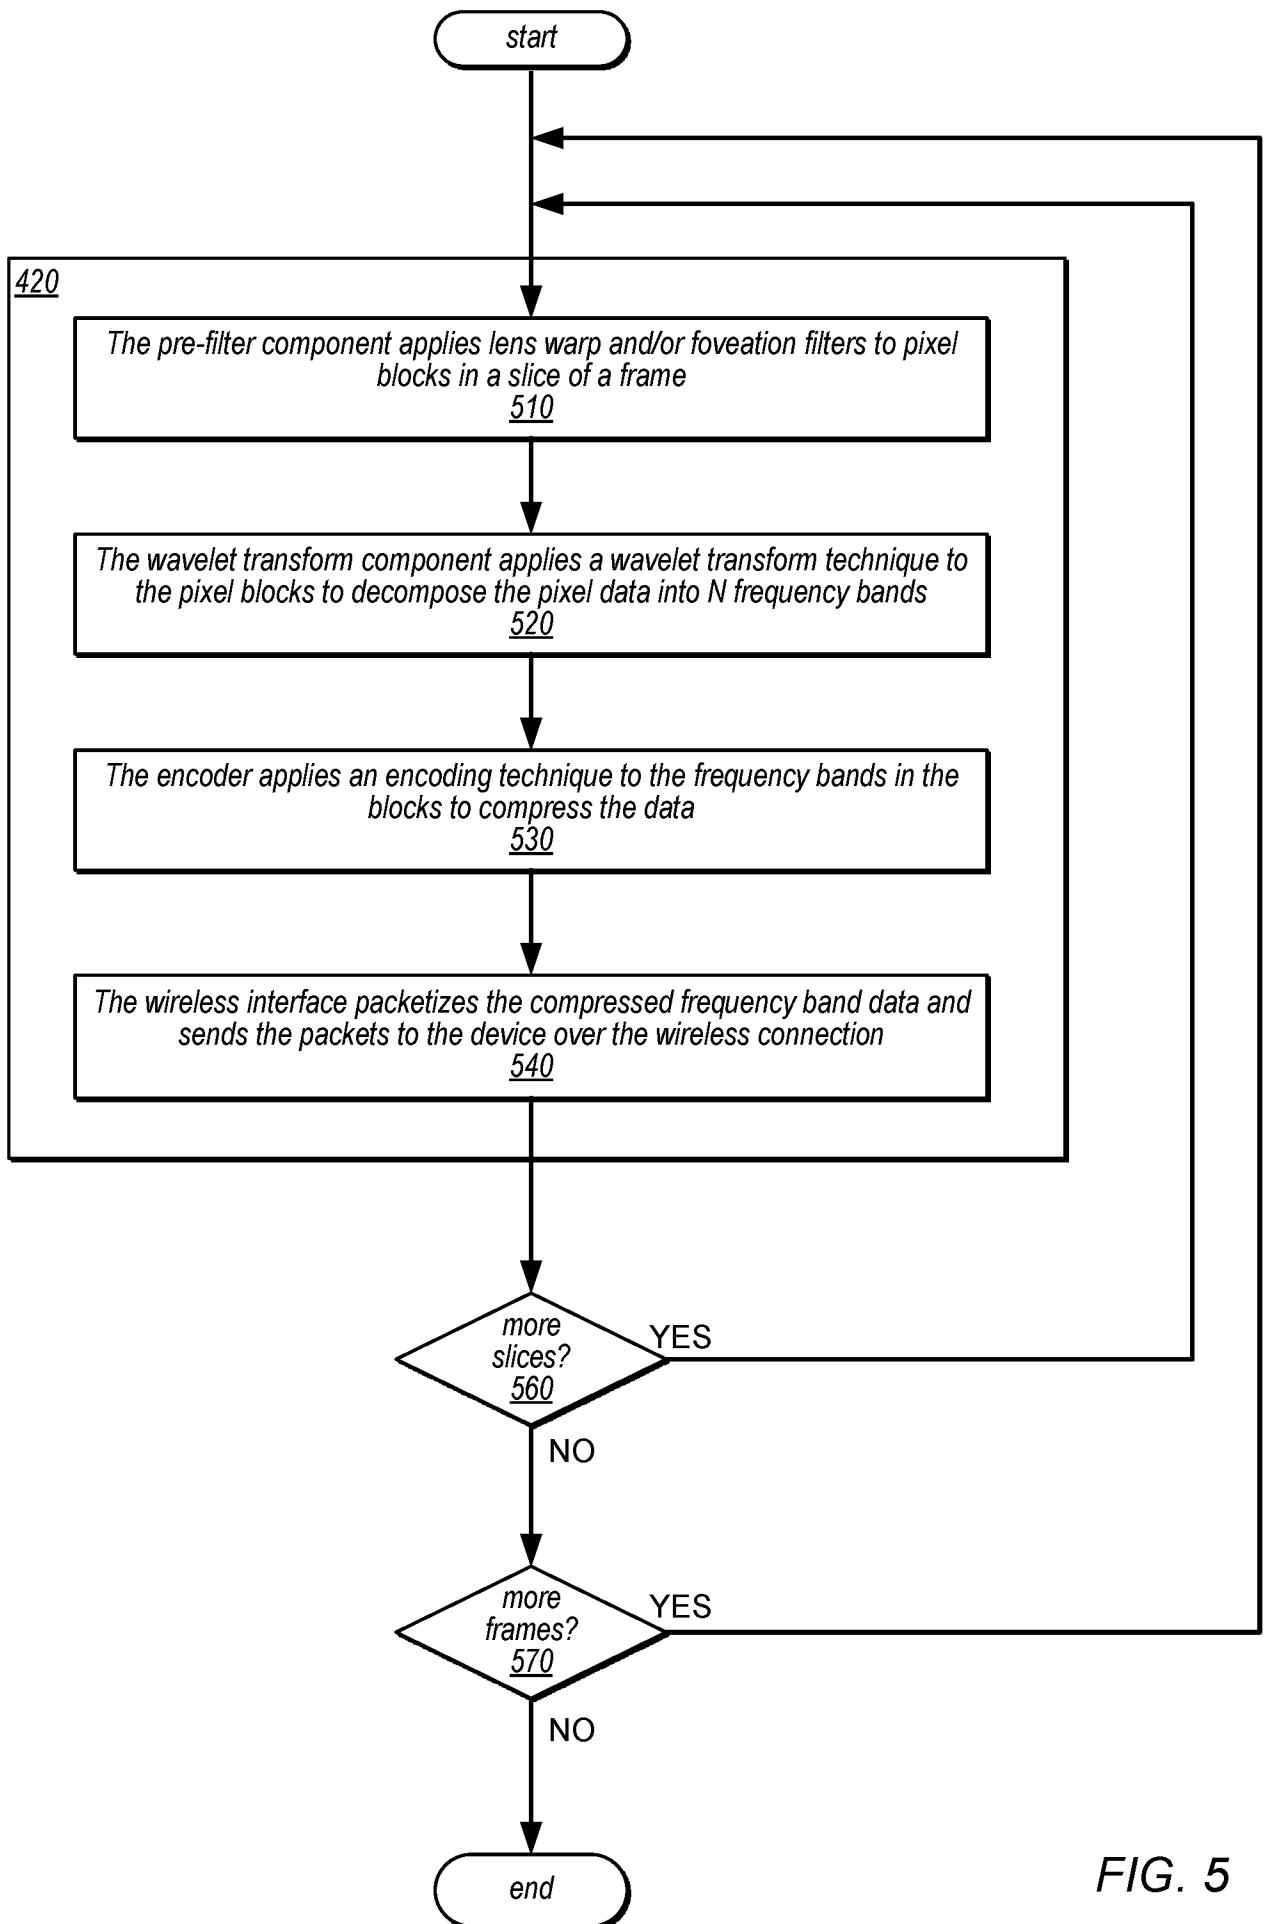


FIG. 5

7 / 9

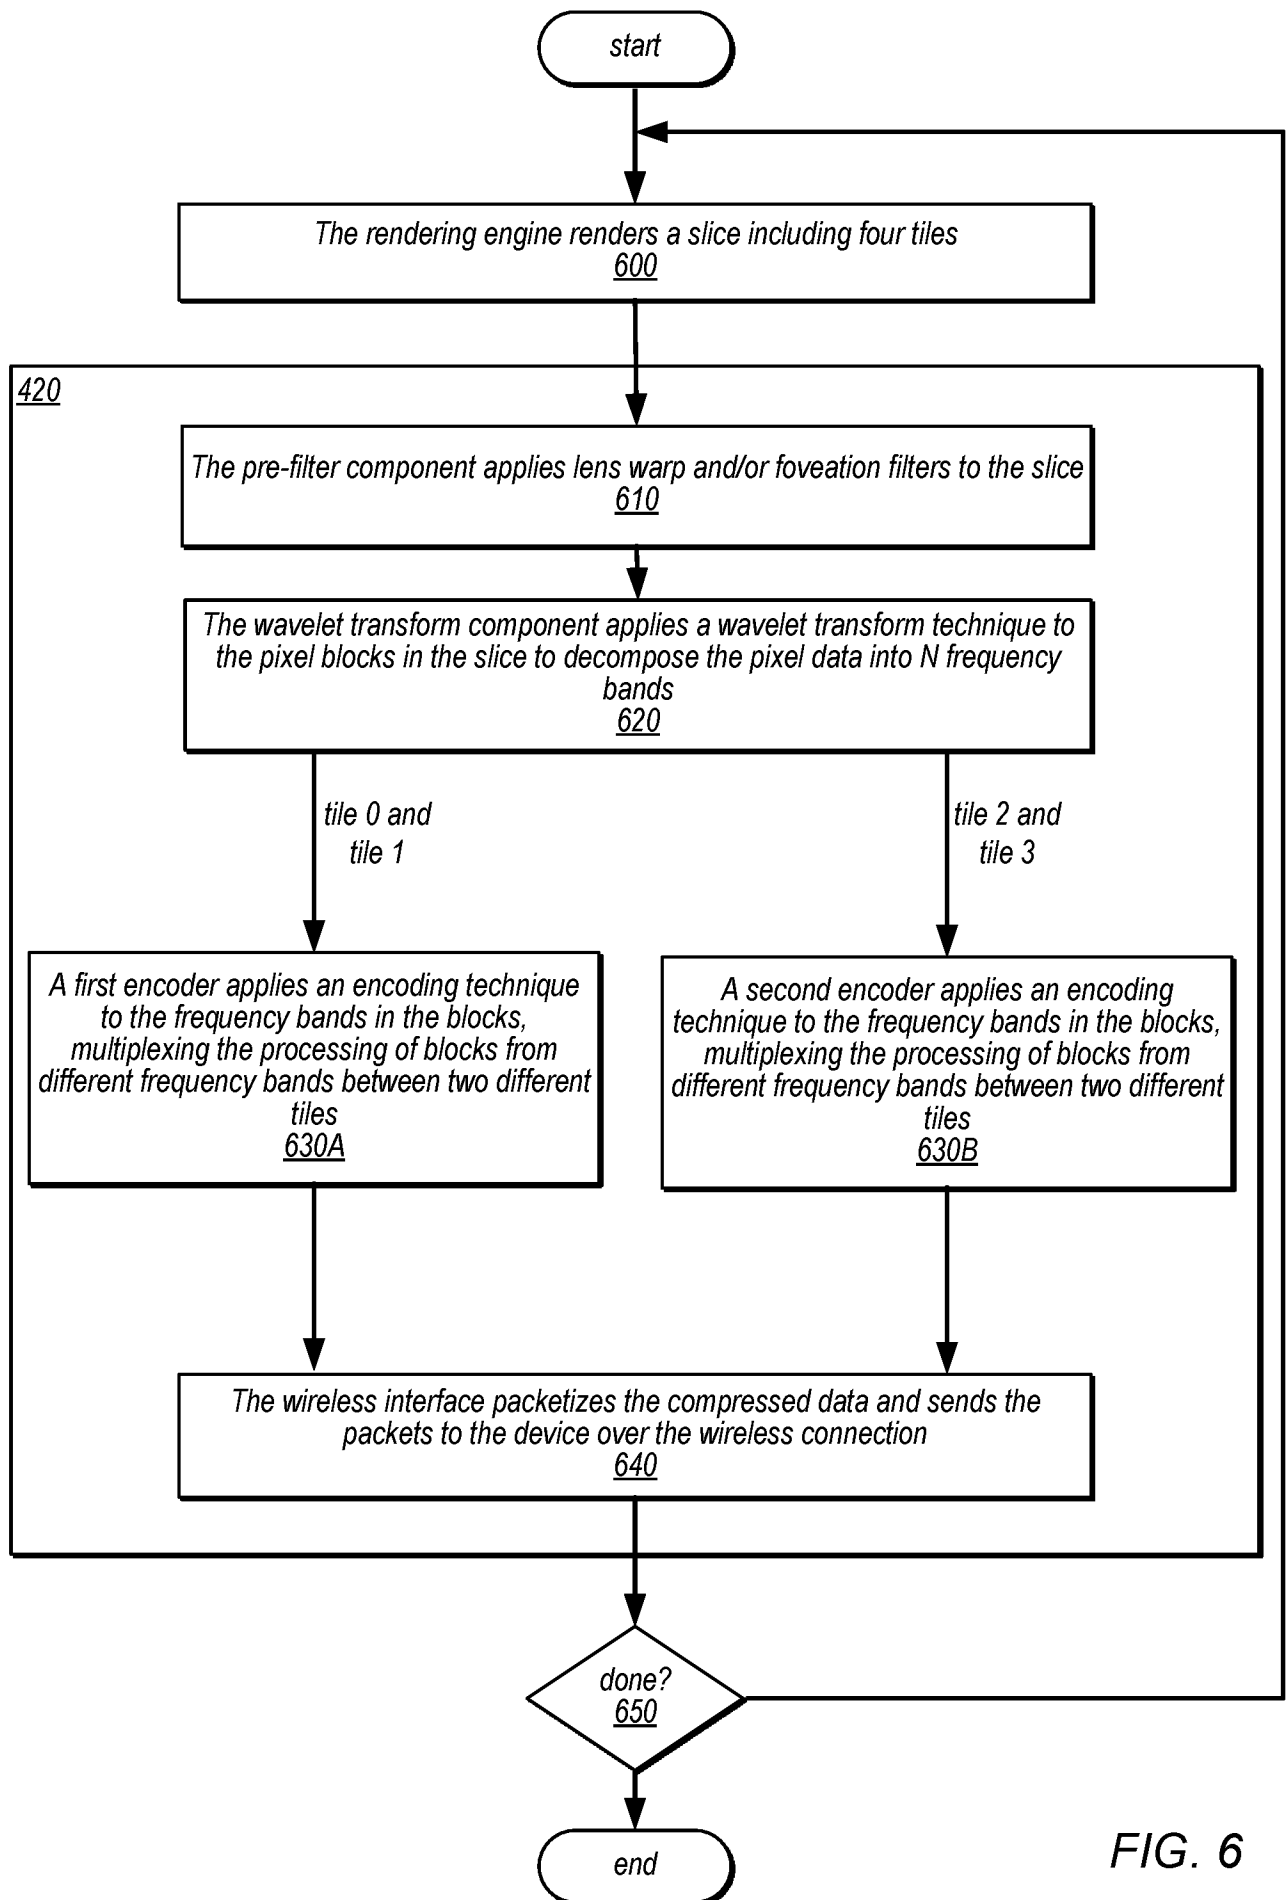


FIG. 6




FIG. 7

(WORLD)

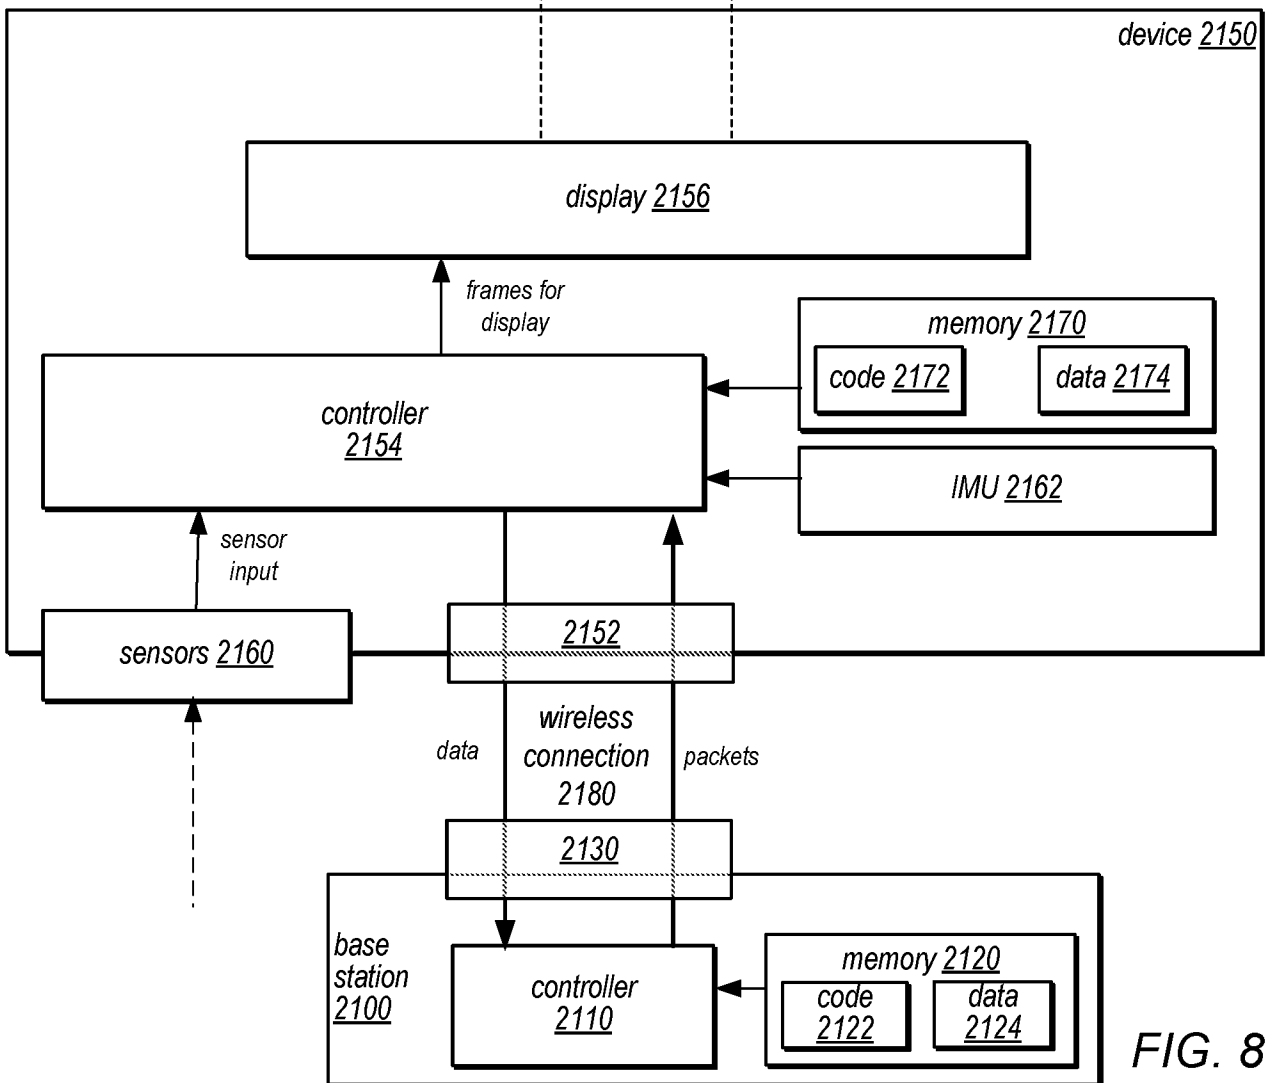



FIG. 8

INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2019/039595

A. CLASSIFICATION OF SUBJECT MATTER  
INV. H04N19/167 H04N19/80 H04N19/86 H04N19/63  
ADD.  
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED  
Minimum documentation searched (classification system followed by classification symbols)  
H04N  
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 1 301 021 A2 (EASTMAN KODAK CO [US]) 9 April 2003 (2003-04-09)	1,5-9, 11,12, 14-18,20
Y	paragraph [0012] - paragraph [0018]; figure 1 paragraph [0038] - paragraph [0053]; figures 4-5	2-4,10, 13,19
Y	----- WO 2005/032137 A1 (UNIV CALIFORNIA [US]; MOLINO ANDREA [IT] ET AL.) 7 April 2005 (2005-04-07) page 4, line 25 - page 5, line 13; figures 1-2 ----- -/--	2

Further documents are listed in the continuation of Box C.

See patent family annex.

\* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier application or patent but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
- "&" document member of the same patent family

Date of the actual completion of the international search  10 October 2019	Date of mailing of the international search report  23/10/2019
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer  Georgiou, Georgia

## INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2019/039595

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2017/301065 A1 (ADSUMILLI BALINEEDU CHOWDARY [US] ET AL) 19 October 2017 (2017-10-19) paragraph [0008] paragraph [0080] paragraph [0102] paragraph [0104] paragraph [0122]; figure 6A -----	2,3,13, 19
Y	US 2017/287112 A1 (STAFFORD JEFFREY ROGER [US] ET AL) 5 October 2017 (2017-10-05) paragraph [0021] - paragraph [0024]; figure 2 -----	4,13
Y	ANJUL PATNEY ET AL: "Towards foveated rendering for gaze-tracked virtual reality", ACM TRANSACTIONS ON GRAPHICS, ACM, 2 PENN PLAZA, SUITE 701NEW YORKNY10121-0701USA, vol. 35, no. 6, 11 November 2016 (2016-11-11), pages 1-12, XP058306343, ISSN: 0730-0301, DOI: 10.1145/2980179.2980246 the whole document -----	4,13
Y	A.; Langelaar Hanjalic ET AL: "Chapter 5 Noise reduction by coring" In: "ADVANCES IN IMAGE COMMUNICATION.", 1 January 2000 (2000-01-01), ELSEVIER, AMSTERDAM, NL, XP055629019, ISSN: 0928-1479 vol. 8, pages 93-121, DOI: 10.1016/S0928-1479(00)80007-8, section [5.1 Introduction] section [5.3 5.3 Coring image sequences in wavelet and subband domains] -----	10
A	US 2009/189981 A1 (SIANN JON [US] ET AL) 30 July 2009 (2009-07-30) paragraph [0096] - paragraph [0106]; figure 1A paragraph [0241] -----	1-20

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2019/039595

Patent document cited in search report	Publication date	Patent family member(s)	Publication date	
EP 1301021	A2	09-04-2003	EP 1301021 A2	09-04-2003
			JP 2003202850 A	18-07-2003
			US 2003067476 A1	10-04-2003
			US 2004207632 A1	21-10-2004
			US 2004207635 A1	21-10-2004
-----				
WO 2005032137	A1	07-04-2005	CN 1883204 A	20-12-2006
			EP 1665796 A1	07-06-2006
			US 2005105812 A1	19-05-2005
			WO 2005032137 A1	07-04-2005
-----				
US 2017301065	A1	19-10-2017	NONE	
-----				
US 2017287112	A1	05-10-2017	NONE	
-----				
US 2009189981	A1	30-07-2009	EP 2238758 A2	13-10-2010
			US 2009189981 A1	30-07-2009
			US 2011096168 A1	28-04-2011
			US 2016173833 A1	16-06-2016
			WO 2009094591 A2	30-07-2009
-----				