



(19) **United States**

(12) **Patent Application Publication**
Abuan et al.

(10) **Pub. No.: US 2011/0249077 A1**

(43) **Pub. Date: Oct. 13, 2011**

(54) **VIDEO CONFERENCE NETWORK
MANAGEMENT FOR A MOBILE DEVICE**

Publication Classification

(51) **Int. Cl.**
H04N 7/14 (2006.01)
(52) **U.S. Cl.** **348/14.02; 348/E07.078**
(57) **ABSTRACT**

(76) Inventors: **Joe S. Abuan**, San Jose, CA (US);
James O. Normile, Los Altos, CA
(US); **Hsi-Jung Wu**, Sunnyvale, CA
(US); **Yan Yang**, San Jose, CA
(US); **Hyeonkuk Jeong**, Saratoga,
CA (US); **Xiaosong Zhou**,
Campbell, CA (US); **Roberto
Garcia, JR.**, Sunnyvale, CA (US)


Some embodiments provide a method for managing a video conference between a first device and a second device. The method identifies a first ceiling bit rate for transmitting video conference data to the second device through the communication channel. The method identifies a current bit rate that is less than the first ceiling bit rate. The method receives networking data regarding the communication channel from the second device. The method determines, from the received network data, that the communication channel will sustain an increase in the current bit rate. The method increments the current bit rate. The method iteratively performs the receiving, determining, and incrementing operations until a determination is made that the communication channel will not sustain the increase in the current bit rate.

(21) Appl. No.: **12/794,774**

(22) Filed: **Jun. 6, 2010**

Related U.S. Application Data

(60) Provisional application No. 61/321,871, filed on Apr. 7, 2010.






Figure 2




Figure 1




Figure 3





Figure 4



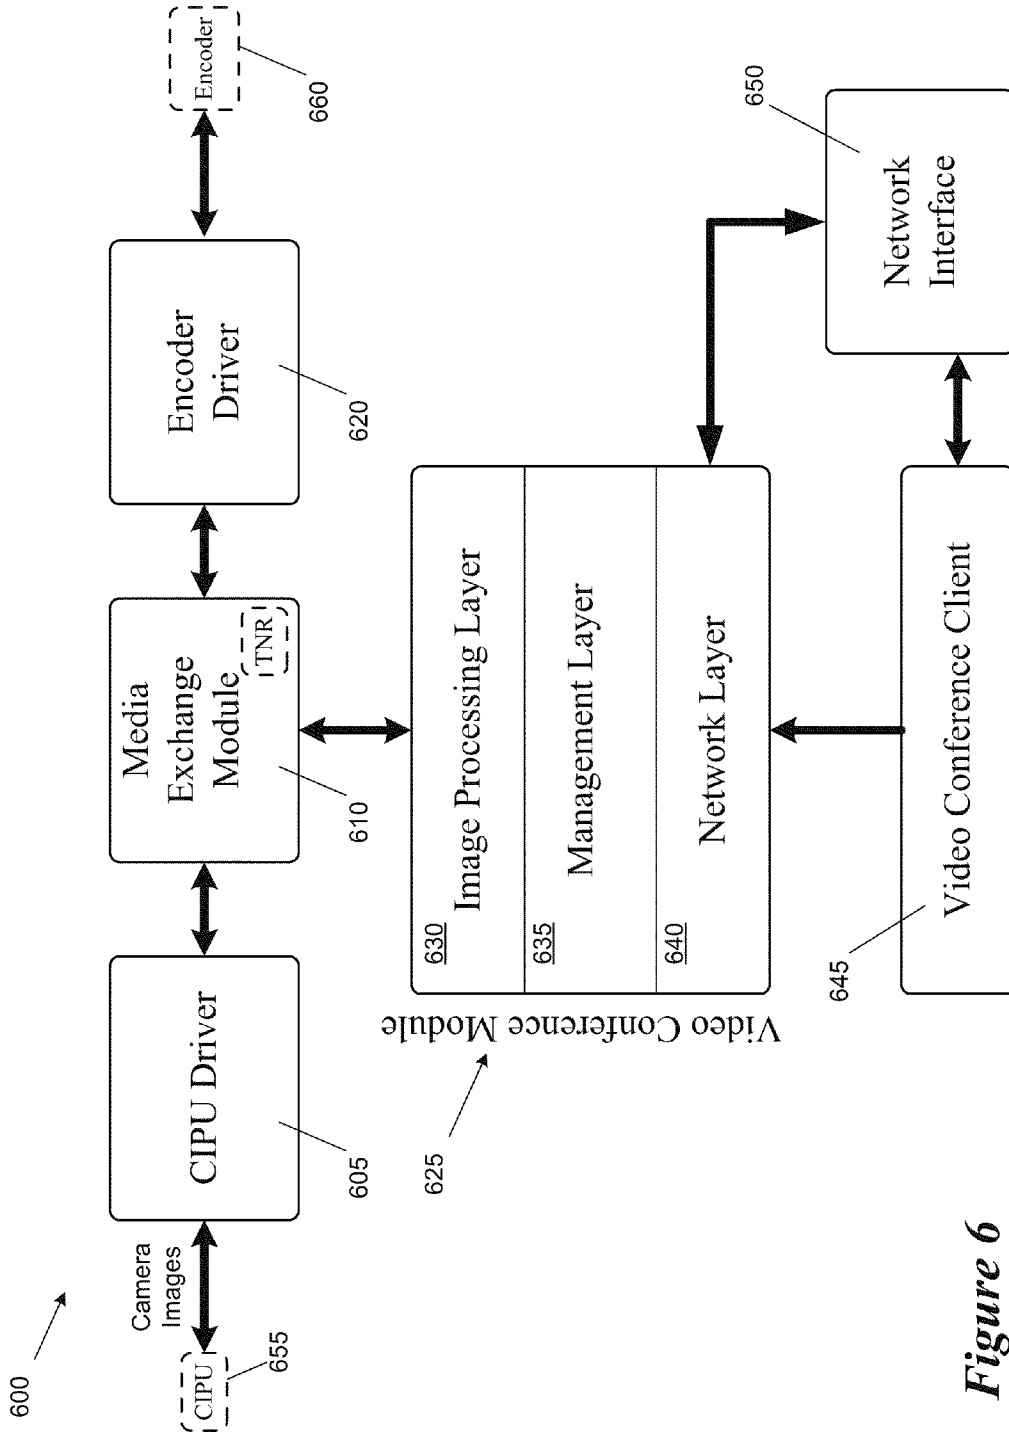


Figure 6

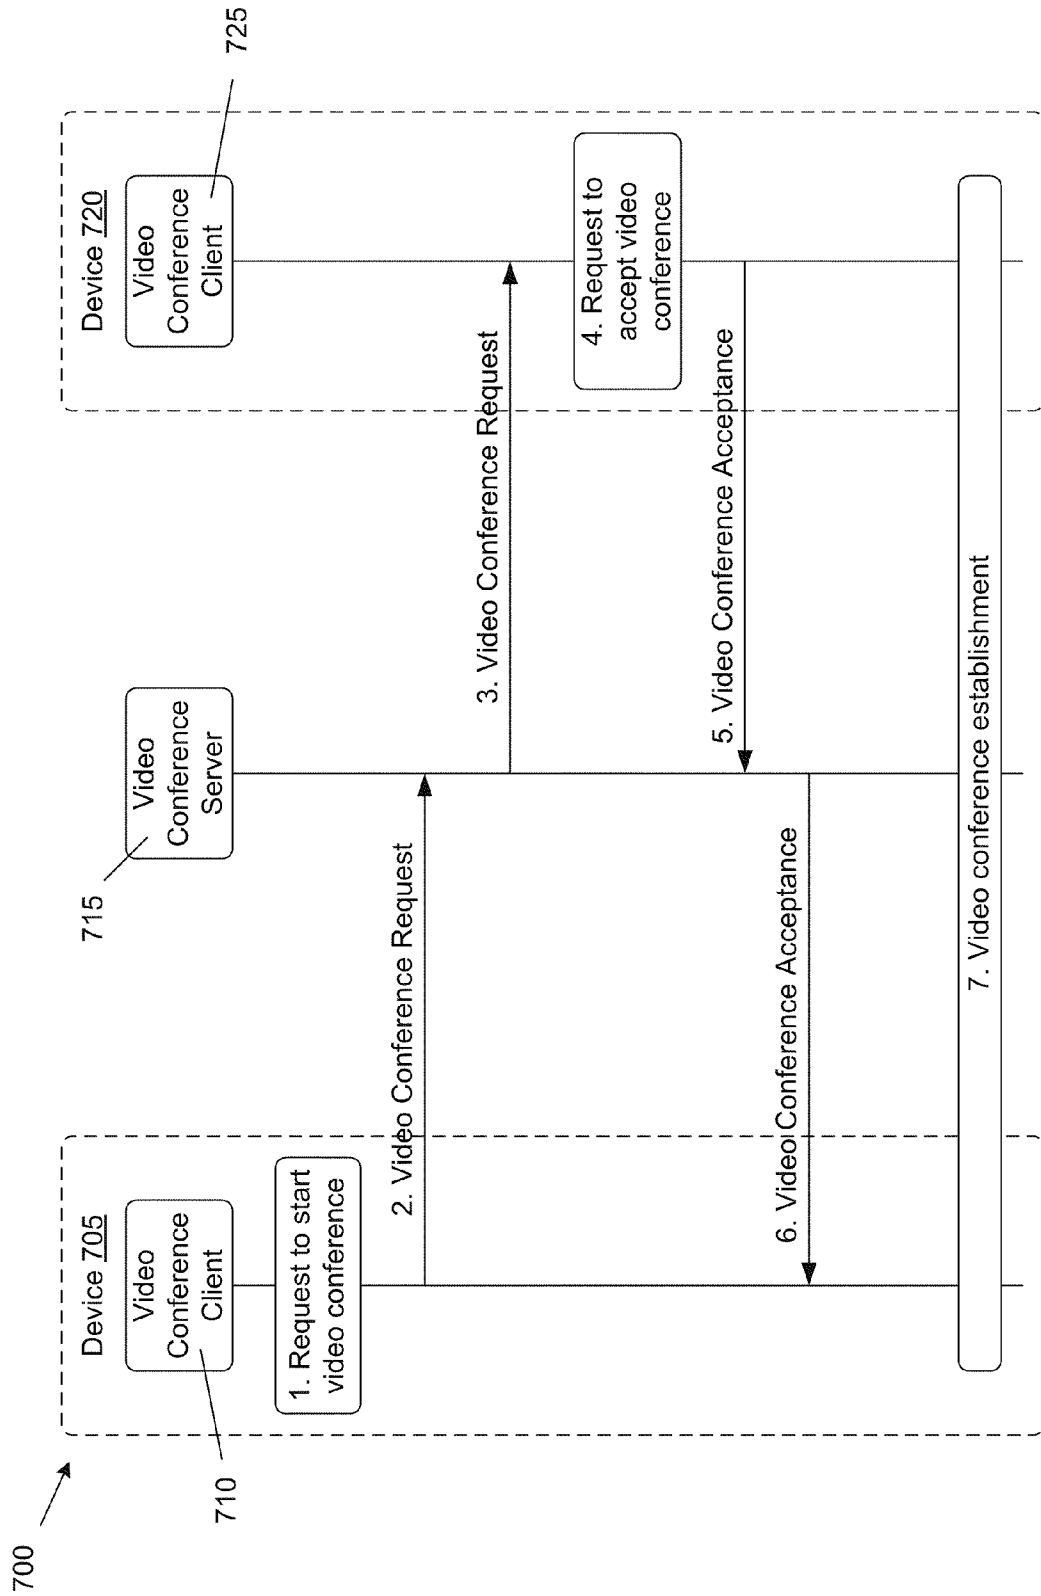


Figure 7




Figure 8




Figure 9




Figure 10




Figure 11

1200 ↘

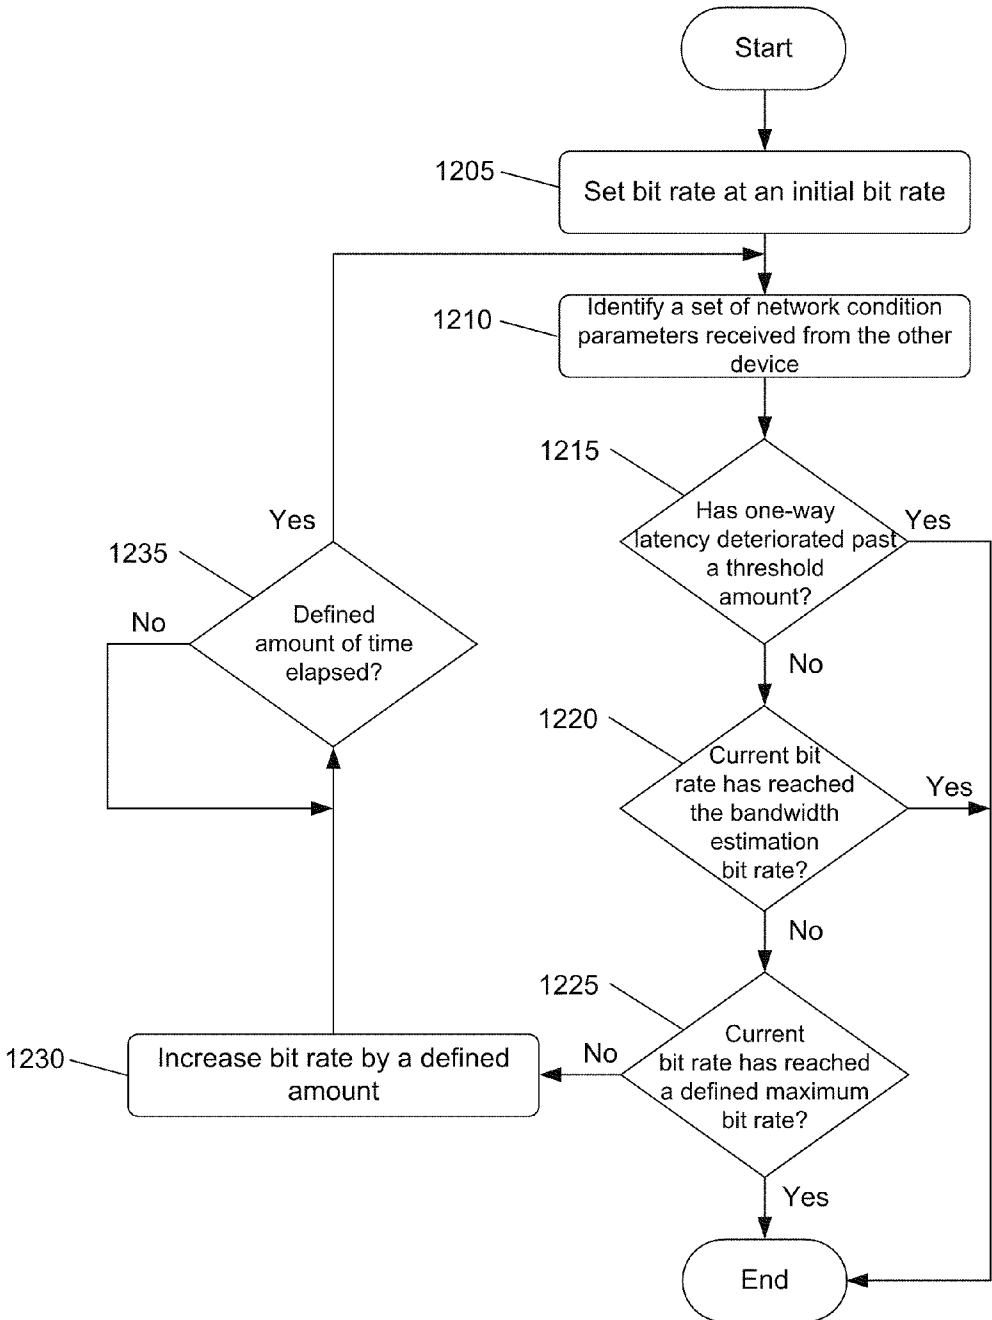


Figure 12

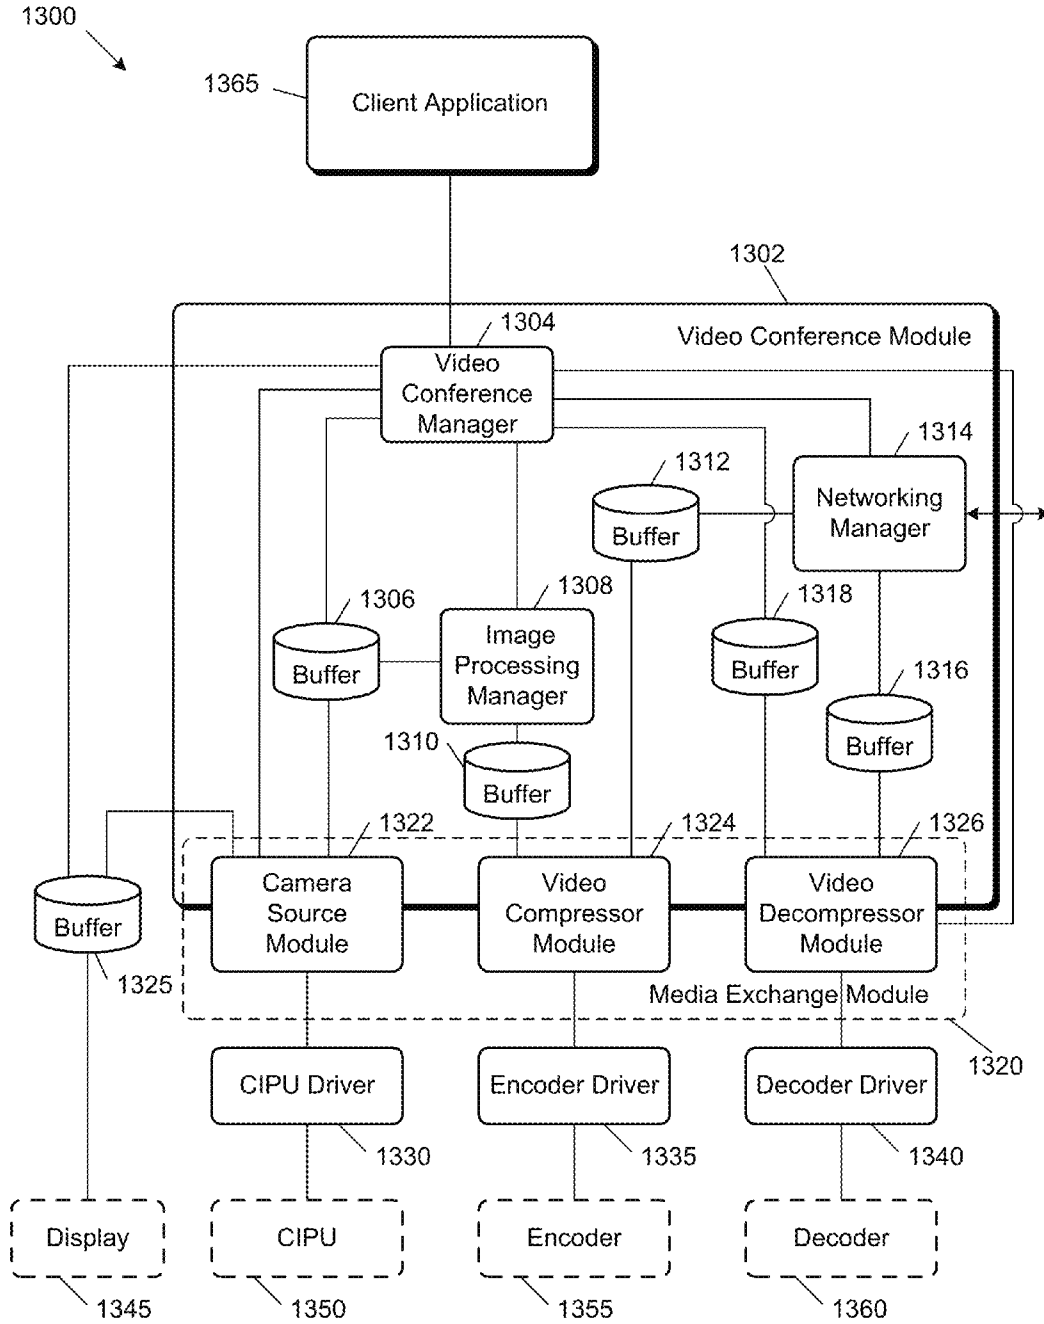


Figure 13

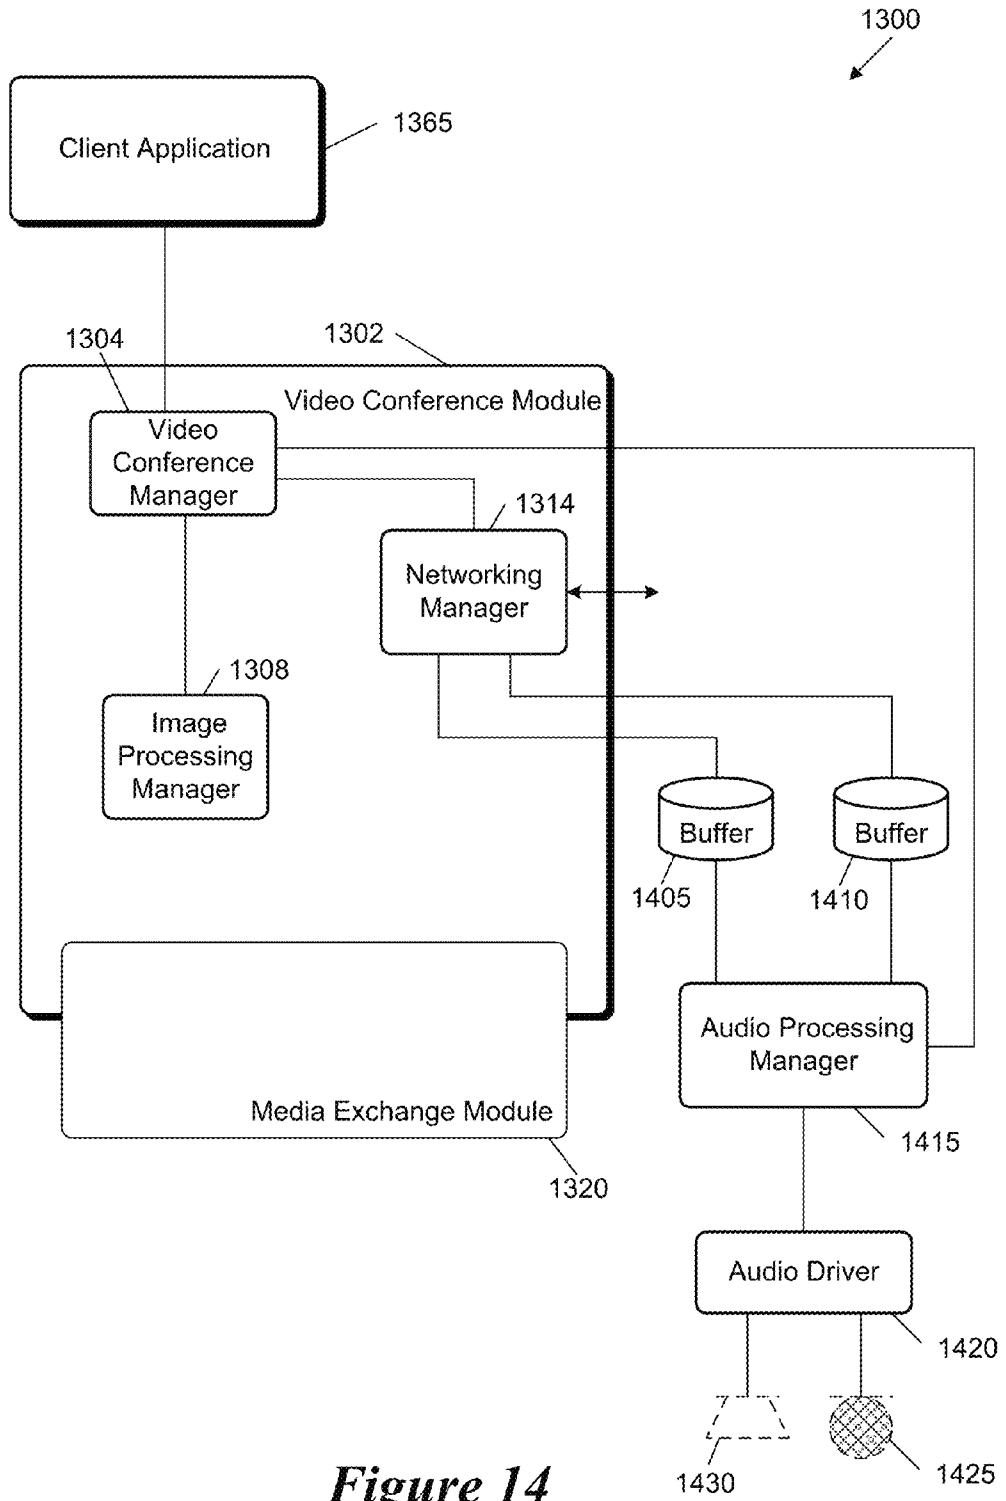


Figure 14

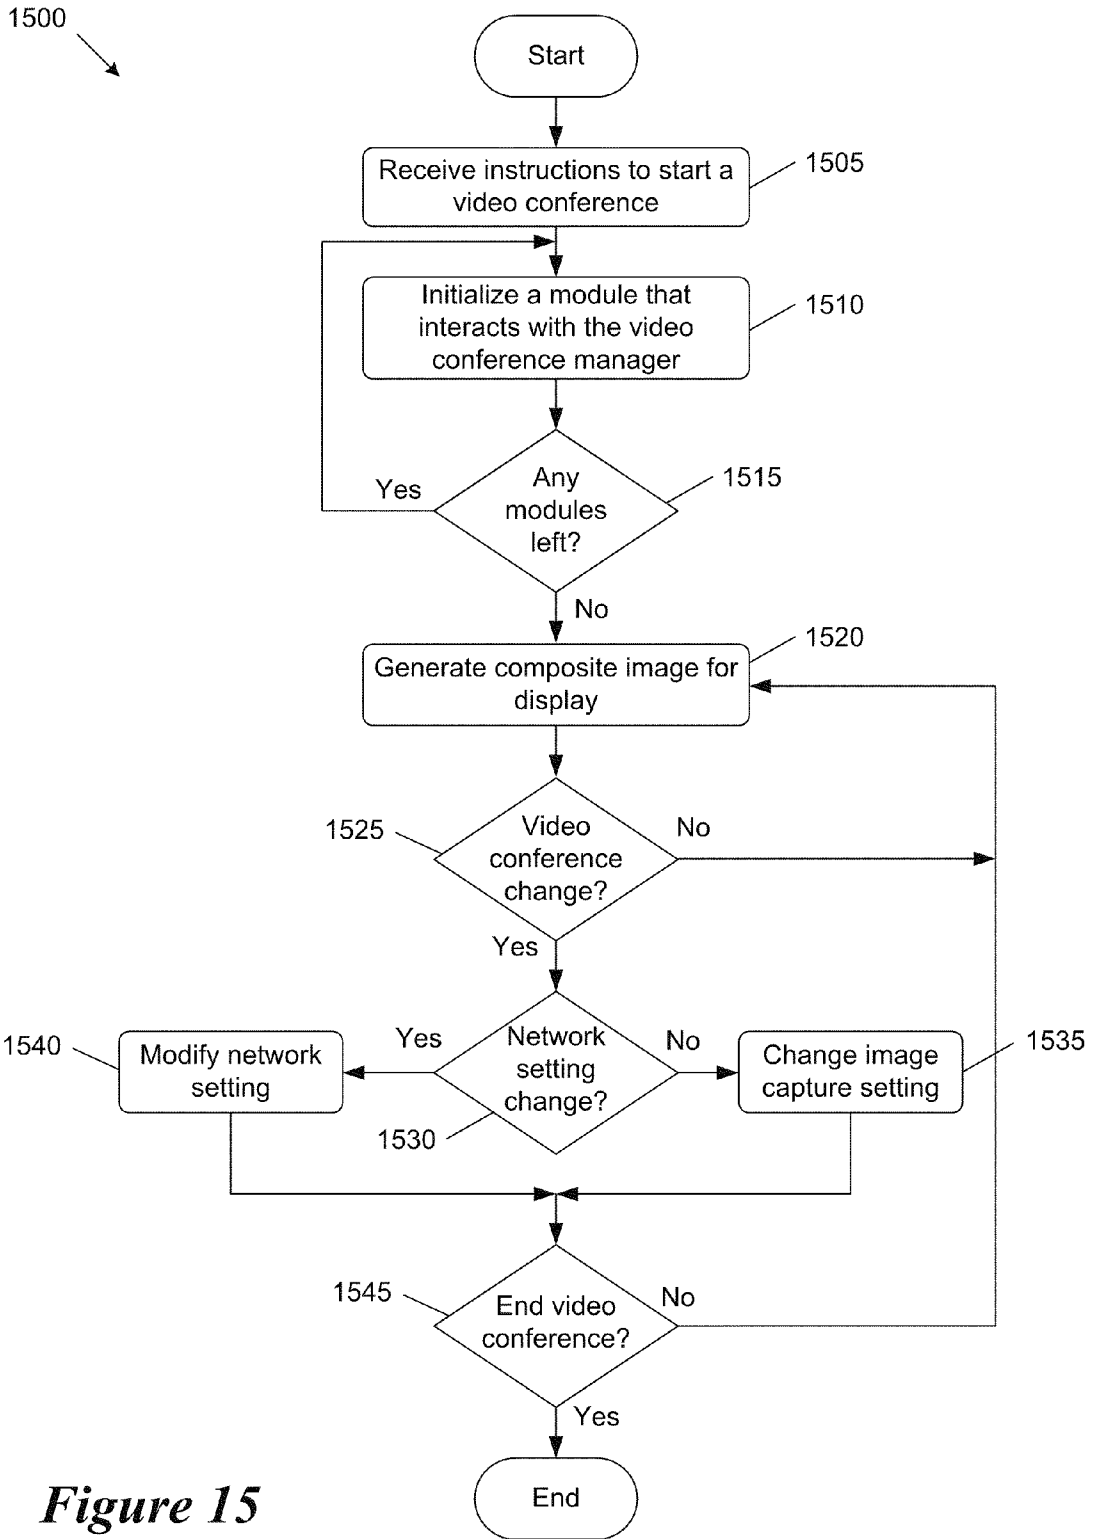


Figure 15

1600

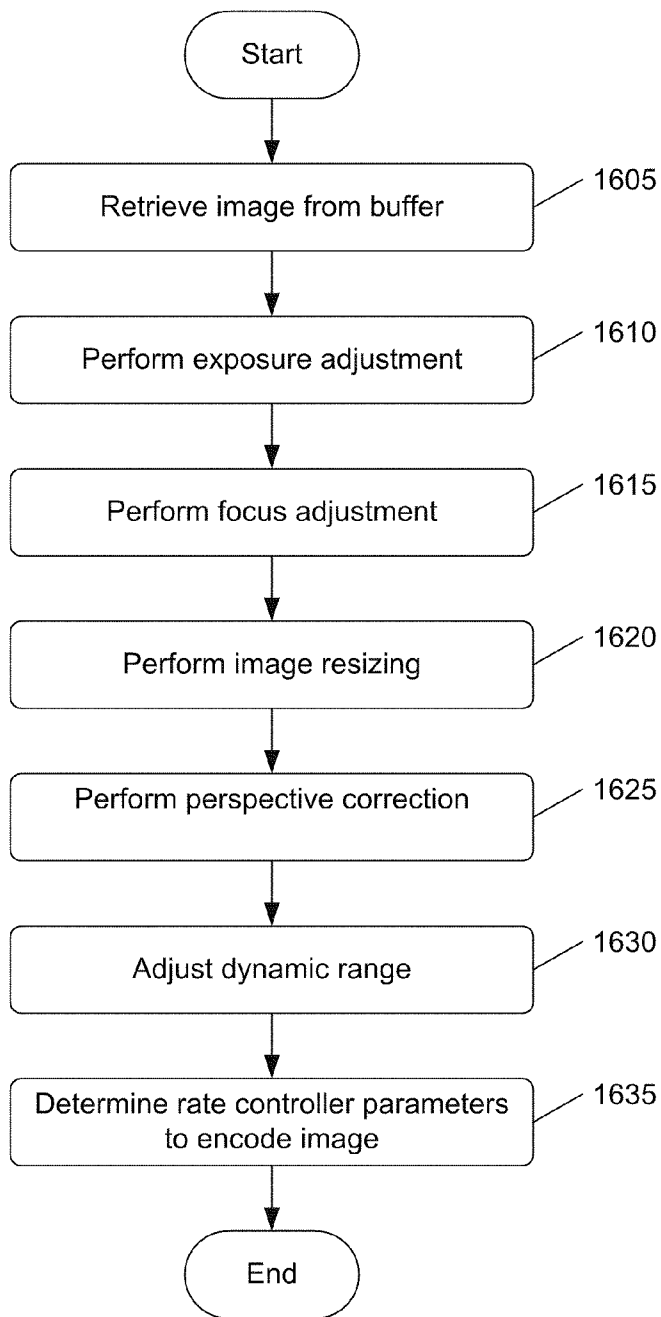


Figure 16




Figure 17

1800 ↘

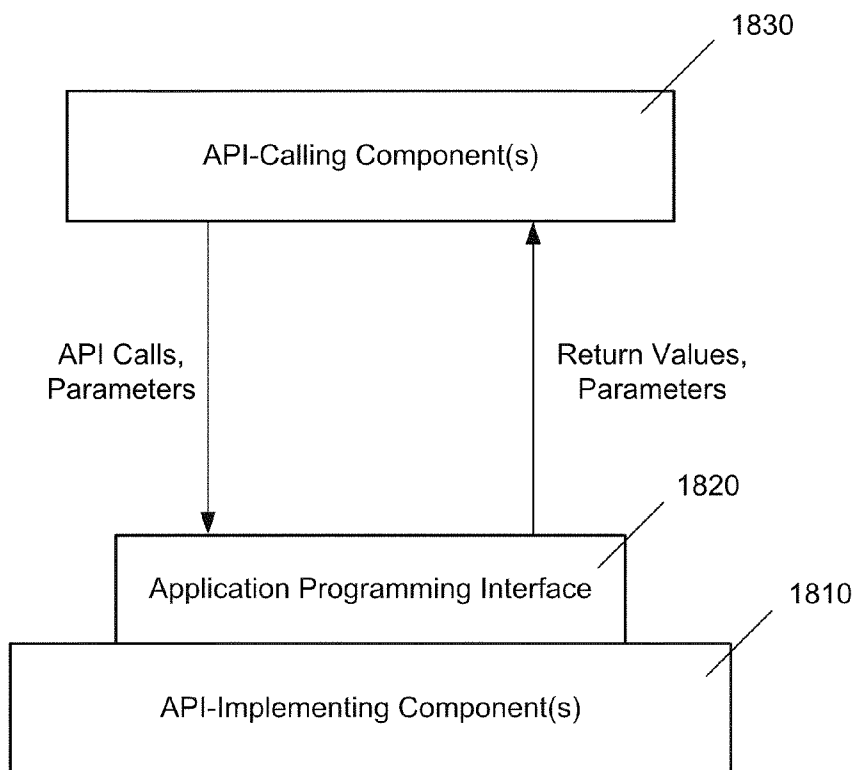


Figure 18

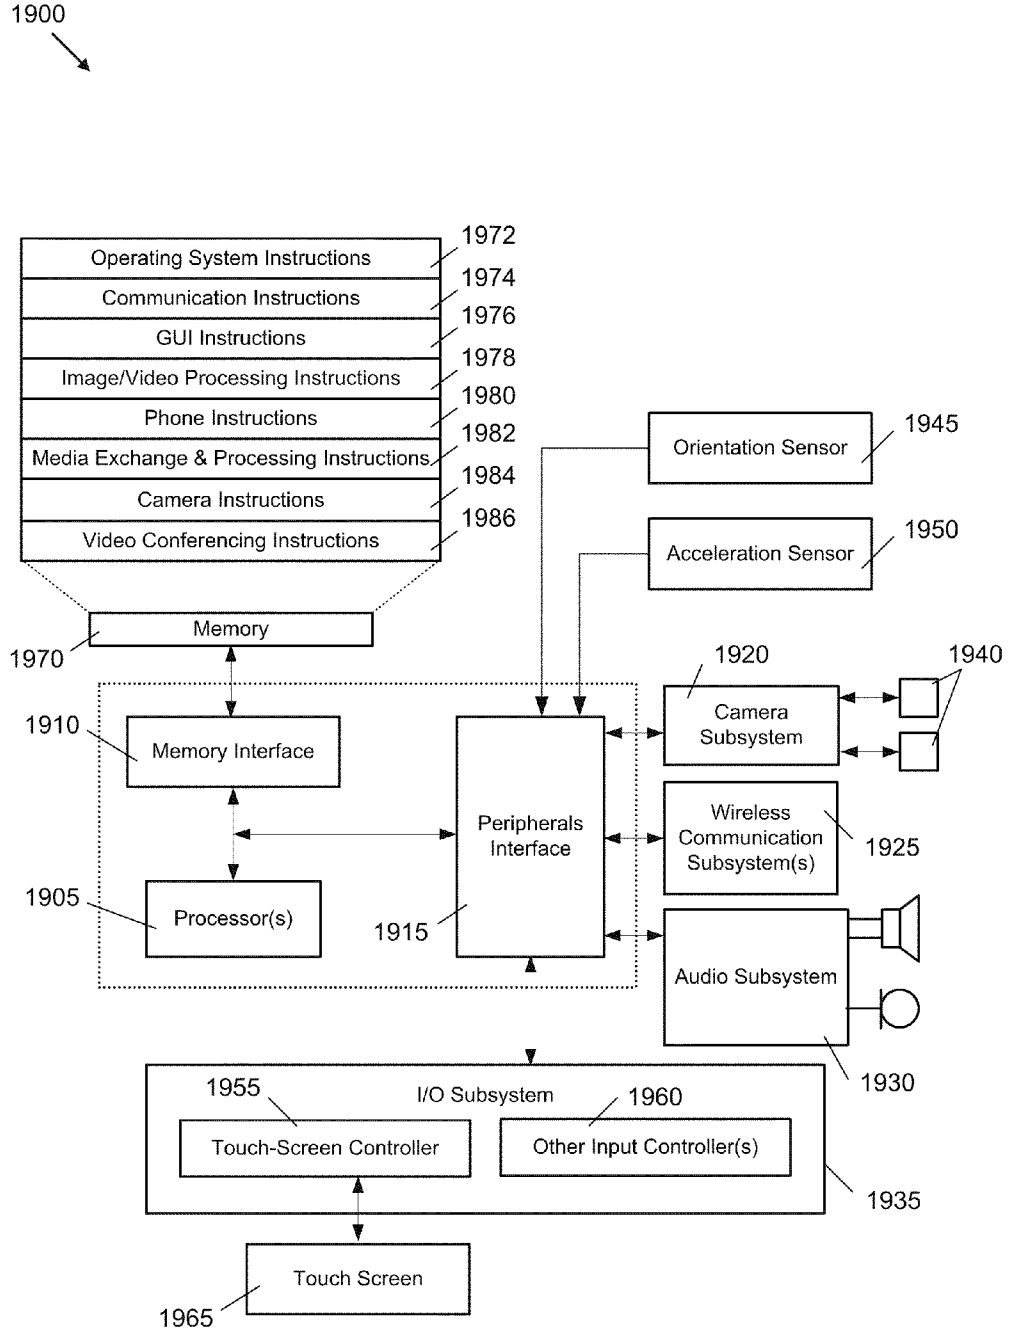


Figure 19

2000
↓




Figure 20




Figure 21




Figure 22

VIDEO CONFERENCE NETWORK MANAGEMENT FOR A MOBILE DEVICE

BACKGROUND

[0001] Many of today's portable devices, such as smartphones, provide video capture functionality. A user of the portable device can capture both still images and video through a camera on the phone. However, to transmit captured video to another party, the user must generally either send the video directly to the other party or upload the video to another location (e.g., an Internet video hosting site) after the video is done being captured. Unfortunately, this does not allow the other party to view the live video stream as it is captured by the portable device.

[0002] In addition, standard portable devices are only equipped with one camera, and processing information from this one camera is difficult enough. An ideal device would have multiple cameras and could send out live video that is a composition of video from at least two cameras. This is an especially difficult problem in light of the limited resources available for portable devices, both in terms of the device processing multiple captured video streams and a network to which the device is connected handling the transmission of the live video streams.

BRIEF SUMMARY

[0003] Some embodiments of the invention provide a mobile device with two cameras that can take pictures and videos. The mobile device of some embodiments has a display screen for displaying the captured picture images and video images. It also includes a storage for storing the captured images for later transmission to another device. The device further has a network interface that allows the device to transmit the captured images to one or more devices during a real-time communication session between the users of the devices. The device also includes an encoder that it can use to encode the captured images for local storage or for transmission to another device. The mobile device further includes a decoder that allows the device to decode images captured by another device during a real-time communication session or to decode images stored locally.

[0004] One example of a real-time communication session that involves the transmission of the captured video images is a video conference. In some embodiments, the mobile device can only transmit one camera's captured video images at any given time during a video conference. In other embodiments, however, the mobile device can transmit captured video images from both of its cameras simultaneously during a video conference or other real-time communication session.

[0005] During a video conference with another device, the mobile device of some embodiments can transmit other types of content along with the video captured by one or both of its cameras. One example of such other content includes low or high resolution picture images that are captured by one of the device's cameras, while the device's other camera is capturing a video that is used in the video conference. Other examples of such other content include (1) files and other content stored on the device, (2) the screen display of the device (i.e., the content that is displayed on the device's screen), (3) content received from another device during a video conference or other real-time communication session, etc.

[0006] Some embodiments provide a novel method for managing the video conference between a first mobile device and a second device. The method identifies a first ceiling bit rate for transmitting video conference data to the second device through the communication channel. The method identifies a current bit rate that is less than the first ceiling bit rate. The method receives networking data regarding the communication channel from the second device. The method determines, from the received network data, that the communication channel will sustain an increase in the current bit rate. The method increments the current bit rate. The method iteratively performs the receiving, determining, and incrementing operations until a determination is made that the communication channel will not sustain the increase in the current bit rate.

[0007] The preceding Summary is intended to serve as a brief introduction to some embodiments of the invention. It is not meant to be an introduction or overview of all inventive subject matter disclosed in this document. The Detailed Description that follows and the Drawings that are referred to in the Detailed Description will further describe the embodiments described in the Summary as well as other embodiments. Accordingly, to understand all the embodiments described by this document, a full review of the Summary, Detailed Description and the Drawings is needed.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] The novel features of the invention are set forth in the appended claims. However, for purpose of explanation, several embodiments of the invention are set forth in the following figures.

[0009] FIG. 1 illustrates a composite display of some embodiments.

[0010] FIG. 2 illustrates another composite display of some embodiments.

[0011] FIG. 3 conceptually illustrates a software architecture for a video processing and encoding module of a dual camera mobile device of some embodiments.

[0012] FIG. 4 conceptually illustrates a captured image processing unit of some embodiments.

[0013] FIG. 5 conceptually illustrates examples of different frame rates based on different vertical blanking intervals (VBIs).

[0014] FIG. 6 conceptually illustrates a software architecture for a video conferencing and processing module of a dual camera mobile device of some embodiments.

[0015] FIG. 7 conceptually illustrates an example video conference request messaging sequence of some embodiments.

[0016] FIG. 8 illustrates a user interface of some embodiments for a video conference setup operation.

[0017] FIG. 9 illustrates a user interface of some embodiments for accepting an invitation to a video conference.

[0018] FIG. 10 illustrates another user interface of some embodiments for accepting an invitation to a video conference.

[0019] FIG. 11 illustrates another user interface of some embodiments for a video conference setup operation.

[0020] FIG. 12 conceptually illustrates a process of some embodiments for setting a bit rate for a video conference.

[0021] FIG. 13 conceptually illustrates another software architecture for a video conferencing and processing module of a dual camera mobile device of some embodiments.

[0022] FIG. 14 conceptually illustrates another software architecture for a dual camera mobile device of some embodiments.

[0023] FIG. 15 conceptually illustrates a process performed by a video conference manager of some embodiments such as that illustrated in FIG. 13.

[0024] FIG. 16 conceptually illustrates a process performed by an image processing manager of some embodiments such as that illustrated in FIG. 6.

[0025] FIG. 17 conceptually illustrates a software architecture for a networking manager of some embodiments such as that illustrated in FIG. 13.

[0026] FIG. 18 conceptually illustrates an application programming interface (API) architecture of some embodiments.

[0027] FIG. 19 illustrates an architecture for a dual camera mobile computing device of some embodiments.

[0028] FIG. 20 conceptually illustrates a touch input/output (I/O) device of some embodiments.

[0029] FIG. 21 conceptually illustrates an example communication system of some embodiments.

[0030] FIG. 22 conceptually illustrates another example communication system of some embodiments.

DETAILED DESCRIPTION

[0031] In the following description, numerous details are set forth for purpose of explanation. However, one of ordinary skill in the art will realize that the invention may be practiced without the use of these specific details. In other instances, well-known structures and devices are shown in block diagram form in order not to obscure the description of the invention with unnecessary detail.

[0032] Some embodiments of the invention provide a mobile device with two cameras that can take pictures and videos. Examples of mobile devices include mobile phones, smartphones, personal digital assistants (PDAs), laptops, tablet personal computers, or any other type of mobile computing device. As used in this document, pictures refer to still picture images that are taken by the camera one at a time in a single-picture mode, or several at a time in a fast-action mode. Video, on the other hand, refers to a sequence of video images that are captured by a camera at a particular rate, which is often referred to as a frame rate. Typical frame rates for capturing video are 25 frames per second (fps), 30 fps, and 60 fps. The cameras of the mobile device of some embodiments can capture video images (i.e., video frames) at these and other frame rates.

[0033] The mobile device of some embodiments (1) can display the captured picture images and video images, (2) can store the captured images for later transmission to another device, (3) can transmit the captured images to one or more devices during a real-time communication session between the users of the devices, and (4) can encode the captured images for local storage or for transmission to another device.

[0034] One example of a real-time communication session that involves the transmission of the captured video images is a video conference. In some embodiments, the mobile device can only transmit one camera's captured video images at any given time during a video conference. In other embodiments, however, the mobile device can transmit captured video images from both of its cameras simultaneously during a video conference or other real-time communication session.

[0035] The mobile devices of some embodiments generate composite displays that include simultaneous display of mul-

tiple videos captured by multiple cameras of one or more devices. In some cases, the composite displays place the videos in adjacent display areas (e.g., in adjacent windows). FIG. 1 illustrates one such example of a composite display 100 that includes two adjacent display areas 105 and 110 that simultaneously display two videos captured by two cameras of one device or captured by two cameras of two different devices that are in a video conference.

[0036] In other cases, the composite display is a PIP display that includes at least two display areas that show two different videos, where one of the display areas is a background main display area and the other is a foreground inset display area that overlaps the background main display area. FIG. 2 illustrates one such example of a composite PIP display 200. This composite PIP display 200 includes a background main display area 205 and a foreground inset display area 210 that overlaps the background main display area. The two display areas 205 and 210 simultaneously display two videos captured by two cameras of one device, or captured by two cameras of two different devices that are in a video conference. While the example composite PIP displays illustrated and discussed in this document are similar to the composite PIP display 200, which shows the entire foreground inset display area 210 within the background main display area 205, other composite PIP displays that have the foreground inset display area 210 overlapping, but not entirely inside, the background main display area 205 are possible.

[0037] In addition to transmitting video content during a video conference with another device, the mobile device of some embodiments can transmit other types of content along with the conference's video content. One example of such other content includes low or high resolution picture images that are captured by one of the device's cameras, while the device's other camera is capturing a video that is used in the video conference. Other examples of such other content include (1) files and other content stored on the device, (2) the screen display of the device (i.e., the content that is displayed on the device's screen), (3) content received from another device during a video conference or other real-time communication session, etc.

[0038] The mobile devices of some embodiments employ novel in-conference adjustment techniques for making adjustments during a video conference. For instance, while transmitting only one camera's captured video during a video conference, the mobile device of some embodiments can dynamically switch to transmitting the video captured by its other camera. In such situations, the mobile device of some embodiments notifies any other device participating in the video conference of this switch so that this other device can provide a smooth transition on its end between the videos captured by the two cameras.

[0039] In some embodiments, the request to switch cameras not only can originate on the "local" device that switches between its cameras during the video conference, but also can originate from the other "remote" device that is receiving the video captured by the local device. Moreover, allowing one device to direct another device to switch cameras is just one example of a remote control capability of the devices of some embodiments. Examples of other operations that can be directed to a device remotely in some embodiments include exposure adjustment operations (e.g., auto-exposure), focus adjustment operations (e.g., auto-focus), etc. Another example of a novel in-conference adjustment that can be specified locally or remotely is the identification of a region of

interest (ROI) in a captured video, and the use of this ROI identification to modify the behavior of the capturing camera, to modify the image processing operation of the device with the capturing camera, or to modify the encoding operation of the device with the capturing camera.

[0040] Yet another example of a novel in-conference adjustment of some embodiments involves real-time modifications of composite video displays that are generated by the devices. Specifically, in some embodiments, the real-time modifications of the composite video displays involve moving one or more of the display areas within a composite display in response to a user's selection and movement of the display areas. Some embodiments also rotate the composite display during a video conference, when the screen of the device that provides this composite display rotates. Also, the mobile device of some embodiments allow the user of the device to flip the order of videos in a PIP display (i.e., to make the video in the foreground inset display appear in the background main display, while making the video in the background main display appear in the foreground inset display).

[0041] Several more detailed embodiments are described below. Section I provides a description of the video processing architecture of some embodiments. Section II then describes the captured image processing unit of some embodiments. In some embodiments, this unit is the component of the device that is responsible for processing raw images captured by the cameras of the device.

[0042] Next, Section III describes the video conferencing architecture of some embodiments. This section also describes the video conference module of some embodiments, as well as several manners for setting up a single camera video conference. Section IV describes the hardware architecture of the dual camera device of some embodiments. Lastly, United States patent application Ser. No. _____, entitled "Establishing a Video Conference During a Phone Call," filed concurrently with the present application, with the Attorney Docket No. APLE.P0212, describes several additional embodiments relating to some of the features described above, such as remote control of mobile devices, camera switch operations, in-conference adjustments, etc. This United States patent application entitled "Establishing a Video Conference During a Phone Call" is incorporated herein by reference.

I. Video Capture and Processing

[0043] FIG. 3 conceptually illustrates a video processing and encoding module 300 of a dual camera mobile device of some embodiments. In some embodiments, the module 300 processes images and encodes videos that are captured by the cameras of the dual camera mobile device. As shown in FIG. 3, this module 300 includes a captured image processing unit (CIPU) driver 305, a media exchange module 310, an encoder driver 320, and a video processing module 325.

[0044] In some embodiments, the media exchange module 310 allows programs on the device that are consumers and producers of media content to exchange media content and instructions regarding the processing of the media content. In the video processing and encoding module 300, the media exchange module 310 of some embodiments routes instructions and media content between the video processing module 325 and the CIPU driver 305, and between the video processing module 325 and the encoder driver 320. To facilitate the routing of such instructions and media content, the media exchange module 310 of some embodiments provides

a set of application programming interfaces (APIs) for the consumers and producers of media content to use. In some of such embodiments, the media exchange module 310 is a set of one or more frameworks that is part of an operating system running on the dual camera mobile device. One example of such a media exchange module 310 is the Core Media framework provided by Apple Inc.

[0045] The video processing module 325 performs image processing on the images and/or the videos captured by the cameras of the device. Examples of such operations include exposure adjustment operations, focus adjustment operations, perspective correction, dynamic range adjustment, image resizing, image compositing, etc. In some embodiments, some image processing operations can also be performed by the media exchange module 310. For instance, as shown in FIG. 3, the media exchange module 310 of some embodiments performs a temporal noise reduction (TNR) operation (e.g., by TNR 315) that reduces noise in video images captured by the cameras of the device. Further examples of such image processing operations of the video processing module 325 and the media exchange module 310 will be provided below.

[0046] Through the media exchange module 310, the video processing module 325 interfaces with the CIPU driver 305 and the encoder driver 320, as mentioned above. The CIPU driver 305 serves as a communication interface between a captured image processing unit (CIPU) 330 and the media exchange module 310. As further described below, the CIPU 330 is the component of the dual camera device that is responsible for processing images captured during image capture or video capture operations of the device's cameras. From the video processing module 325 through the media exchange module 310, the CIPU driver 305 receives requests for images and/or videos from one or both of the device's cameras. The CIPU driver 305 relays such requests to the CIPU 330, and in response receives the requested images and/or videos from the CIPU 330, which the CIPU driver 305 then sends to the video processing module 325 through the media exchange module 310. Through the CIPU driver 305 and the media exchange module 310, the video processing module 325 of some embodiments also sends instructions to the CIPU 330 in order to modify some of its operations (e.g., to modify a camera's frame rate, exposure adjustment operation, focus adjustment operation, etc.).

[0047] The encoder driver 320 serves as a communication interface between the media exchange module 310 and an encoder hardware 335 (e.g., an encoder chip, an encoding component on a system on chip, etc.). In some embodiments, the encoder driver 320 receives images and requests to encode the images from the video processing module 325 through the media exchange module 310. The encoder driver 320 sends the images to be encoded to the encoder 335, which then performs picture encoding or video encoding on the images. When the encoder driver 320 receives encoded images from the encoder 335, the encoder driver 320 sends the encoded images back to the video processing module 325 through the media exchange module 310.

[0048] In some embodiments, the video processing module 325 can perform different operations on the encoded images that it receives from the encoder. Examples of such operations include storing the encoded images in a storage of the device, transmitting the encoded images in a video conference through a network interface of the device, etc.

[0049] In some embodiments, some or all of the modules of the video processing and encoding module **300** are implemented as part of an operating system. For example, some embodiments implement all four components **305**, **310**, **320**, and **325** of this module **300** as part of the operating system of the device. Other embodiments implement the media exchange module **310**, the CIPU driver **305**, and the encoder driver **320** as part of the operating system of the device, while having the video processing module **325** as an application that runs on the operating system. Still, other implementations of the module **300** are possible.

[0050] The operation of the video processing and encoding module **300** during a video capture session will now be described. To start a video capture session, the video processing module **325** initializes several components that are needed for the video capture session. In some embodiments, these components include (1) the CIPU **330**, (2) a scaling and compositing module (not shown) of the video processing module **325**, (3) an image processing module (not shown) of the video processing module **325**, and (4) the encoder **335**. Also, the video processing module **325** of some embodiments initializes a network manager (not shown) when it is participating in a video conference.

[0051] Through the media exchange module **310** and the CIPU driver **305**, the video processing module sends its initialization request to the CIPU **330**, in order to have one or both of the cameras of the device start video capturing. In some embodiments, this request specifies a particular frame rate, exposure level, and scaling size for each camera that needs to capture a video. In response to this request, the CIPU **330** starts to return video images from the requested cameras at the specified rate(s), exposure level(s), and scaling size(s). These video images are returned to the video processing module **325** through the CIPU driver **305** and the media exchange module **310**, which, as mentioned above, performs TNR operations on the video images before supplying them to the video processing module **325**. At the video processing module **325**, the video images are stored in a buffer (not shown) for additional image processing.

[0052] The image processing module of the video processing module **325** retrieves the video images stored in the buffer for additional video processing. The scaling and compositing module then retrieves the processed video images in order to scale them if necessary for real time display on the display screen of the device. In some embodiments, this module creates composite images from the images captured by two cameras of the device or from images captured by the camera (s) of the device along with the camera(s) of another device during a video conference in order to provide a real-time display of the captured video images on the device or to create a composite video image for encoding.

[0053] The processed and/or composited video images are supplied to the encoder **335** through the encoder driver **320** and the media exchange module **310**. The encoder **335** then encodes the video images. The encoded images are then returned to the video processing module **325** (again through the encoder driver **320** and the media exchange module **310**) for storage on the device or for transmission during a video conference. When the device is participating in a video conference, the network manager (that was initialized by the video processing module **325**) then retrieves these encoded

images, packetizes them and transmits them to one or more other devices through a network interface (not shown) of the device.

II. Captured Image Processing

[0054] The images captured by cameras of the dual camera mobile device of some embodiments are raw, unprocessed images. These images require conversion to a particular color space before the images can be used for other operations such as transmitting the images to another device (e.g., during a video conference), storing the images, or displaying the images. In addition, the images captured by the cameras may need to be processed to correct errors and/or distortions and to adjust the images' color, size, etc. Accordingly, some embodiments perform several processing operations on the images before storing, transmitting, and displaying such images. Part of the processing of such images is performed by the CIPU **330**.

[0055] One example of such a CIPU is illustrated in FIG. 4. Specifically, this figure conceptually illustrates a captured image processing unit (CIPU) **400** of some embodiments. This CIPU **400** includes a single processing pipeline **485** that either processes images from only one of the device's cameras at a time, or processes images from both of the device's cameras simultaneously in a time-division multiplex fashion (i.e., in a time interleaved manner). The CIPU **400**'s processing pipeline **485** can be configured differently to address differing characteristics and/or operational settings of the different cameras. Examples of different camera characteristics in some embodiments include different resolutions, noise sensors, lens types (fixed or zoom lens), etc. Also, examples of different operational settings under which the device can operate the cameras in some embodiments include image resolution size, frame rate, zoom level, exposure level, etc.

[0056] As shown in FIG. 4, the CIPU **400** includes a sensor module **415**, a line/frame buffer **417**, a bad pixel correction (BPC) module **420**, a lens shading (LS) module **425**, a demosaicing module **430**, a white balance (WB) module **435**, a gamma module **440**, a color space conversion (CSC) module **445**, a hue, saturation, and contrast (HSC) module **450**, a scaler module **455**, a filter module **460**, a statistics engine **465**, two sets of registers **470**, and a controller module **475**. In some embodiments, all of the modules of the CIPU **400** are implemented in hardware (e.g., an ASIC, FPGA, a SOC with a microcontroller, etc.), while in other embodiments, some or all of the modules of the CIPU **400** are implemented in software.

[0057] As shown in FIG. 4, the sensor module **415** communicatively couples to two pixel arrays **410a** and **410b** and two sets of sensors **405a** and **405b** of two cameras of the device. In some embodiments, this communicative coupling is facilitated through each camera sensor's mobile industry processor interface (MIPI).

[0058] Through this communicative coupling, the sensor module **415** can forward instructions to the cameras to control various aspects of each camera's operations such as its power level, zoom level, focus, exposure level, etc. In some embodiments, each camera has four operational power modes. In the first operational power mode, the camera is powered off. For the second operational power mode, the camera is powered on, but it is not yet configured. In the third operational power mode, the camera is powered on, the camera's sensor is configured, and the camera sensor's pixels are collecting photons and converting the collected photons to digital values. How-

ever, the camera sensor is not yet sending images to the sensor module 415. Finally, in the fourth operational power mode, the camera is in the same operational power mode as the third power mode except the camera is now sending images to the sensor module 415.

[0059] During the operation of the device, the cameras may switch from one operational power mode to another any number of times. When switching operational power modes, some embodiments require the cameras to switch operational power modes in the order described above. Therefore, in those embodiments, a camera in the first operational power mode can only switch to the second operational power mode. When the camera is in the second operational power mode, it can switch to the first operational power mode or to the third operational power mode. Similarly, the camera can switch from the third operational power mode to the second operational power mode or the fourth operational power mode. When the camera is in the fourth operational power mode, it can only switch back to the third operational power mode.

[0060] Moreover, switching from one operational power mode to the next or the previous operational power mode takes a particular amount of time. Thus, switching between two or three operational power modes is slower than switching between one operational power mode. The different operational power modes also consume different amounts of power. For instance, the fourth operational power mode consumes the most amount of power, the third operational power mode consumes more power than the first and second, and the second operational power mode consumes more than the first. In some embodiments, the first operational power mode does not consume any power.

[0061] When a camera is not in the fourth operational power mode capturing images, the camera may be left in one of the other operational power modes. Determining the operational mode in which to leave the unused camera depends on how much power the camera is allowed to consume and how fast the camera may need to respond to a request to start capturing images. For example, a camera configured to operate in the third operational power mode (e.g., standby mode) consumes more power than a camera configured to be in the first operational power mode (i.e., powered off). However, when the camera is instructed to capture images, the camera operating in the third operational power mode can switch to the fourth operational power mode faster than the camera operating in the first operational power mode. As such, the cameras can be configured to operate in the different operational power modes when not capturing images based on different requirements (e.g., response time to a request to capture images, power consumption).

[0062] Through its communicative coupling with each camera, the sensor module 415 can direct one or both sets of camera sensors to start capturing images when the video processing module 325 requests one or both cameras to start capturing images and the sensor module 415 receives this request through the controller module 475, as further described below. Bayer filters are superimposed over each of the camera sensors and thus each camera sensor outputs Bayer pattern images, which are stored in the pixel array associated with each camera sensor. A Bayer pattern image is an image where each pixel only stores one color value: red, blue, or green.

[0063] Through its coupling with the pixel arrays 410a and 410b, the sensor module 415 retrieves raw Bayer pattern images stored in the camera pixel arrays 410a and 410b. By

controlling the rate at which the sensor module 415 retrieves images from a camera's pixel array, the sensor module 415 can control the frame rate of the video images that are being captured by a particular camera. By controlling the rate of its image retrieval, the sensor module 415 can also interleave the fetching of images captured by the different cameras in order to interleave the CIPU processing pipeline 485's image processing of the captured images from the different cameras. The sensor module 415's control of its image retrieval is further described below and in the above-incorporated U.S. patent application Ser. No. _____, entitled "Establishing Video Conference During a Phone Call," with Attorney Docket Number APLE.P0212.

[0064] The sensor module 415 stores image lines (i.e., rows of pixels of an image) in the line/frame buffer 417, which the sensor module 415 retrieves from the pixel arrays 410a and 410b. Each image line in the line/frame buffer 417 is processed through the CIPU processing pipeline 485. As shown in FIG. 4, the CIPU processing pipeline 485 is formed by the BPC module 420, the LS module 425, the demosaicing module 430, the WB module 435, the gamma module 440, the CSC module 445, the HSC module 450, the scaler module 455, and the filter module 460. In some embodiments, the CIPU processing pipeline 485 processes images from the line/frame buffer 417 on a line-by-line (i.e., row-by-row) basis while in other embodiments the CIPU processing pipeline 485 processes entire images from the line/frame buffer 417 on a frame-by-frame basis.

[0065] In the exemplary pipeline illustrated in FIG. 4, the BPC module 420 is the module that retrieves the images from the line/frame buffer 417. This module performs a bad-pixel removal operation that attempts to correct bad pixels in the retrieved images that might have resulted from one or more of the camera sensors being defective (e.g., the defective photo sensors do not sense light at all, sense light incorrectly, etc.). In some embodiments, the BPC module 420 detects bad pixels by comparing a particular pixel in an image with one or more neighboring pixels in the image. If the difference between the value of the particular pixel and the values of the neighboring pixels is greater than a threshold amount, the particular pixel's value is replaced by the average of several neighboring pixel's values that are of the same color (i.e., red, green, and blue) as the particular pixel.

[0066] The operation of the BPC module 420 is in part controlled by the values stored for this module in the two sets of registers 470 of the CIPU 400. Specifically, to process the images captured by the two different cameras of the device, some embodiments configure the CIPU processing pipeline 485 differently for each camera, as mentioned above. The CIPU processing pipeline 485 is configured for the two different cameras by storing two different sets of values in the two different sets of registers 470a (Ra) and 470b (Rb) of the CIPU 400. Each set of registers 470 includes one register (Ra or Rb) for each of the modules 420-460 within the CIPU processing pipeline 485. Each register in each register set stores a set of values that defines one processing pipeline module's operation. Accordingly, as shown in FIG. 4, the register set 470a is for indicating the mode of operation of each processing pipeline module for one camera (camera A) of the dual camera mobile device, while the register set 470b is for indicating the mode of operation of each module for the other camera (camera B) of the dual camera mobile device.

[0067] One example of configuring the CIPU processing pipeline 485 differently for each camera is to configure the

modules of the CIPU processing pipeline **485** to process different sized images. For instance, if the camera sensor **405a** is 640×480 pixels and the camera sensor **405b** is 2048×1536 pixels, the set of registers **470a** is configured to store values that instruct the modules of the CIPU processing pipeline **485** to process 640×480 pixel images and the set of registers **470b** is configured to store values that instruct the modules of the CIPU processing pipeline **485** to process 2048×1536 pixel images.

[0068] In some embodiments, different processing pipeline configurations (i.e., register values) are stored in different profile settings. In some of such embodiments, a user of the mobile device is allowed to select one of the profile settings (e.g., through a user interface displayed on the mobile device) to set the operation of a camera(s). For example, the user may select a profile setting for configuring a camera to capture high resolution video, a profile setting for configuring the same camera to capture low resolution video, or a profile setting for configuring both cameras to capture high resolution still images. Different configurations are possible, which can be stored in many different profile settings. In other of such embodiments, instead of allowing the user to select a profile setting, a profile setting is automatically selected based on which application or activity the user selects. For instance, if the user selects a video conferencing application, a profile that configures both cameras to capture video is automatically selected, if the user selects a photo application, a profile that configures one of the cameras to capture still images is automatically selected, etc.

[0069] After the BPC module **420**, the LS module **425** receives the bad-pixel-corrected images. The LS module **425** performs a lens shading correction operation to correct for image defects that are caused by camera lenses that produce light falloff effects (i.e., light is reduced towards the edges of the camera sensor). Such effects cause images to be unevenly illuminated (e.g., darker at corners and/or edges). To correct these image defects, the LS module **425** of some embodiments estimates a mathematical model of a lens' illumination fall-off. The estimated model is then used to compensate the lens fall-off of the image to evenly illuminate unevenly illuminated portions of the image. For example, if a corner of the image is half the brightness of the center of the image, the LS module **425** of some embodiments multiplies the corner pixels value by two in order to produce an even image.

[0070] The demosaicing module **430** performs a demosaicing operation to generate full color images from images of sampled colors. As noted above, the camera sensors output Bayer pattern images, which are incomplete because each pixel of a Bayer pattern image stores only one color value. The demosaicing module **430** reconstructs a red, green, blue (RGB) image from a Bayer pattern image by interpolating the color values for each set of colors in the Bayer pattern image.

[0071] The WB module **435** performs a white balance operation on the RGB images received from the demosaicing module **430** so that the colors of the content of the images are similar to the colors of such content perceived by the human eye in real life. The WB module **435** adjusts the white balance by adjusting colors of the images to render neutral colors (e.g., gray, white, etc.) correctly. For example, an image of a piece of white paper under an incandescent light may appear yellow whereas the human eye perceives the piece of paper as white. To account for the difference between the color of the images that the sensor captures and what the human eye

perceives, the WB module **435** adjusts the color values of the image so that the captured image properly reflects the colors perceived by the human eye.

[0072] The statistics engine **465** collects image data at various stages of the CIPU processing pipeline **485**. For example, FIG. 4 shows that the statistics engine **465** collects image data after the LS module **425**, the demosaicing module **430**, and the WB module **435**. Different embodiments collect data from any number of different stages of the CIPU processing pipeline **485**. The statistics engine **465** processes the collected data, and, based on the processed data, adjusts the operations of the camera sensors **405a** and **405b** through the controller module **475** and the sensor module **415**. Examples of such operations include exposure and focus. Although FIG. 4 shows the statistics engine **465** controlling the camera sensors **405a** and **405b** through the controller module **475**, other embodiments of the statistics engine **465** control the camera sensors through just the sensor module **415**.

[0073] The processed data can also be used to adjust the operations of various modules of the CIPU **400**. For instance, the statistics engine **465** of some embodiments adjusts the operations of the WB module **435** based on data collected after the WB module **435**. In some of such embodiments, the statistics engine **465** provides an automatic white balance (AWB) function by using the processed data to adjust the white balancing operation of the WB module **435**. Other embodiments can use processed data collected from any number of stages of the CIPU processing pipeline **485** to adjust the operations of any number of modules within the CIPU processing pipeline **485**. Further, the statistics engine **465** can also receive instructions from the controller module **475** to adjust the operations of one or more modules of the CIPU processing pipeline **485**.

[0074] After receiving the images from the WB module **435**, the gamma module **440** performs a gamma correction operation on the image to code and decode luminance or tristimulus values of the camera system. The gamma module **440** of some embodiments corrects gamma by converting a 10-12 bit linear signal into an 8 bit non-linear encoding in order to correct the gamma of the image. Some embodiments correct gamma by using a lookup table.

[0075] The CSC module **445** converts the image received from the gamma module **440** from one color space to another color space. Specifically, the CSC module **445** converts the image from an RGB color space to a luminance and chrominance (YUV) color space. However, other embodiments of the CSC module **445** can convert images from and to any number of color spaces.

[0076] The HSC module **450** may adjust the hue, saturation, contrast, or any combination thereof of the images received from the CSC module **445**. The HSC module **450** may adjust these properties to reduce the noise or enhance the images, for example. For instance, the saturation of images captured by a low-noise camera sensor can be increased to make the images appear more vivid. In contrast, the saturation of images captured by a high-noise camera sensor can be decreased to reduce the color noise of such images.

[0077] After the HSC module **450**, the scaler module **455** may resize images to adjust the pixel resolution of the image or to adjust the data size of the image. The scaler module **455** may also reduce the size of the image in order to fit a smaller display, for example. The scaler module **455** can scale the image a number of different ways. For example, the scaler module **455** can scale images up (i.e., enlarge) and down (i.e.,

shrink). The scaler module 455 can also scale images proportionally or scale images anamorphically.

[0078] The filter module 460 applies one or more filter operations to images received from the scaler module 455 to change one or more attributes of some or all pixels of an image. Examples of filters include a low-pass filter, a high-pass filter, a band-pass filter, a bilateral filter, a Gaussian filter, among other examples. As such, the filter module 460 can apply any number of different filters to the images.

[0079] The controller module 475 of some embodiments is a microcontroller that controls the operation of the CIPU 400. In some embodiments, the controller module 475 controls (1) the operation of the camera sensors (e.g., exposure level) through the sensor module 415, (2) the operation of the CIPU processing pipeline 485, (3) the timing of the CIPU processing pipeline 485 (e.g., when to switch camera sensors, when to switch registers, etc.), and (4) a flash/strobe (not shown), which is part of the dual camera mobile device of some embodiments.

[0080] Some embodiments of the controller module 475 process instructions received from the statistics engine 465 and the CIPU driver 480. In some embodiments, the instructions received from the CIPU driver 480 are instructions from the dual camera mobile device (i.e., received from the local device) while in other embodiments the instructions received from the CIPU driver 480 are instructions from another device (e.g., remote control during a video conference). Based on the processed instructions, the controller module 475 can adjust the operation of the CIPU 400 by programming the values of the registers 470. Moreover, the controller module 475 can dynamically reprogram the values of the registers 470 during the operation of the CIPU 400.

[0081] As shown in FIG. 4, the CIPU 400 includes a number of modules in the CIPU processing pipeline 485. However, one of ordinary skill will realize that the CIPU 400 can be implemented with just a few of the illustrated modules or with additional and different modules. In addition, the processing performed by the different modules can be applied to images in sequences different from the sequence illustrated in FIG. 4.

[0082] An example operation of the CIPU 400 will now be described by reference to FIG. 4. For purposes of explanation, the set of registers Ra is used for processing images captured by camera sensor 405a of the dual camera mobile device and the set of registers Rb is used for processing images captured by camera sensor 405b of the dual camera mobile device. The controller module 475 receives instructions from the CIPU driver 480 to produce images captured by one of the cameras of the dual camera mobile device.

[0083] The controller module 475 then initializes various modules of the CIPU processing pipeline 485 to process images captured by one of the cameras of the dual camera mobile device. In some embodiments, this includes the controller module 475 checking that the correct set of registers of the registers 470 are used. For example, if the CIPU driver 480 instructs the controller module 475 to produce images captured by the camera sensor 405a, the controller module 475 checks that the set of registers Ra is the set of registers from which the modules of the CIPU 400 read. If not, the controller module 475 switches between the sets of registers so that the set of registers Ra is the set that is read by the modules of the CIPU 400.

[0084] For each module in the CIPU processing pipeline 485, the mode of operation is indicated by the values stored in

the set of registers Ra. As previously mentioned, the values in the set of registers 470 can be dynamically reprogrammed during the operation of the CIPU 400. Thus, the processing of one image can differ from the processing of the next image. While the discussion of this example operation of the CIPU 400 describes each module in the CIPU 400 reading values stored in registers to indicate the mode of operation of the modules, in some software-implemented embodiments, parameters are instead passed to the various modules of the CIPU 400.

[0085] In some embodiments, the controller module 475 initializes the sensor module 415 by instructing the sensor module 415 to delay a particular amount of time after retrieving an image from the pixel array 410a. In other words, the controller module 475 instructs the sensor module 415 to retrieve the images from the pixel array 410a at a particular rate.

[0086] Next, the controller module 475 instructs the camera sensor 405a through the sensor module 415 to capture images. In some embodiments, the controller module 475 also provides exposure and other camera operation parameters to the camera sensor 405a. In other embodiments, the camera sensor 405a uses default values for the camera sensor operation parameters. Based on the parameters, the camera sensor 405a captures a raw image, which is stored in the pixel array 410a. The sensor module 415 retrieves the raw image from the pixel array 410a and sends the image to the line/frame buffer 417 for storage before the CIPU processing pipeline 485 processing the image.

[0087] Under certain circumstances, images may be dropped by the line/frame buffer 417. When the camera sensors 405a and/or 405b are capturing images at a high rate, the sensor module 415 may receive and store images in the line/frame buffer 417 faster than the BPC module 420 can retrieve the images from the line/frame buffer 417 (e.g., capturing high frame-rate video), and the line/frame buffer 417 will become full. When this happens, the line/frame buffer 417 of some embodiments drops images (i.e., frames) based on a first in, first out basis. That is, when the line/frame buffer 417 drops an image, the line/frame buffer 417 drops the image that was received before all the other images in the line/frame buffer 417.

[0088] The processing of the image by the CIPU processing pipeline 485 starts by the BPC module 420 retrieving the image from the line/frame buffer 417 to correct any bad pixels in the image. The BPC module 420 then sends the image to the LS module 425 to correct for any uneven illumination in the image. After the illumination of the image is corrected, the LS module 425 sends the image to the demosaicing module 430 where it processes the raw image to generate an RGB image from the raw image. Next, the WB module 435 receives the RGB image from the demosaicing module 430 and adjusts the white balance of the RGB image.

[0089] As noted above, the statistics engine 465 may have collected some data at various points of the CIPU processing pipeline 485. For example, the statistics engine 465 collects data after the LS module 425, the demosaicing module 430, and the WB module 435 as illustrated in FIG. 4. Based on the collected data, the statistics engine 465 may adjust the operation of the camera sensor 405a, the operation of one or more modules in the CIPU processing pipeline 485, or both, in order to adjust the capturing of subsequent images from the camera sensor 405a. For instance, based on the collected data, the statistics engine 465 may determine that the exposure

level of the current image is too low and thus instruct the camera sensor 405a through the sensor module 415 to increase the exposure level for subsequently captured images. Thus, the statistics engine 465 of some embodiments operates as a feedback loop for some processing operations.

[0090] After the WB module 435 adjusts the white balance of the image, it sends the image to the gamma module 440 for gamma correction (e.g., adjusting the gamma curve of the image). The CSC module 445 receives the gamma-corrected image from the gamma module 440 and performs color space conversion. In this example, the CSC module 445 converts the RGB image to a YUV image. In other words, the CSC module 445 converts an image that is represented in an RGB color space to an image that is represented in a YUV color space. The HSC module 450 receives the YUV image from the CSC module 445 and adjusts the hue, saturation, and contrast attributes of various pixels in the image. After the HSC module 450, the scaler module 455 resizes the image (e.g., enlarging or shrinking the image). The filter module 460 applies one or more filters on the image after receiving the image from the scaler module 455. Finally, the filter module 460 sends the processed image to the CIPU driver 480.

[0091] In this example of the operation of the CIPU 400 described above, each module in the CIPU processing pipeline 485 processed the image in some manner. However, other images processed by the CIPU 400 may not require processing by all the modules of the CIPU processing pipeline 485. For example, an image may not require white balance adjustment, gamma correction, scaling, or filtering. As such, the CIPU 400 can process images any number of ways based on a variety of received input such as instructions from the CIPU driver 480 or data collected by the statistic engine 465, for example.

[0092] Different embodiments control the rate at which images are processed (i.e., frame rate) differently. One manner of controlling the frame rate is through manipulation of vertical blanking intervals (VBI). For some embodiments that retrieve image lines for processing images on a line-by-line basis, a VBI is the time difference between retrieving the last line of an image of a video captured by a camera of the dual camera mobile device from a pixel array and retrieving the first line of the next image of the video from the pixel array. In other embodiments, a VBI is the time difference between retrieving one image of a video captured by a camera of the dual camera mobile device from a pixel array and retrieving the next image of the video the pixel array.

[0093] One example where VBI can be used is between the sensor module 415 and the pixel arrays 410a and 410b. For example, some embodiments of the sensor module 415 retrieve images from the pixel arrays 410a and 410b on a line-by-line basis and other embodiments of the sensor module 415 retrieve images from the pixel arrays 410a and 410b on an image-by-image basis. Thus, the frame rate can be controlled by adjusting the VBI of the sensor module 415: increasing the VBI reduces the frame rate and decreasing the VBI increases the frame rate.

[0094] FIG. 5 conceptually illustrates examples of different frame rates 505, 510, and 515 based on different VBIs. Each sequence shows an image, which is captured by one of the cameras of the dual camera mobile device, of a person holding a guitar at various time instances 525-555 along timeline 520. In addition, the time between each time instance 525-555 is the same and will be referred to as one time unit. For purposes of explanation, FIG. 5 will now be described by

reference to the sensor module 415 and the pixel array 410a of FIG. 4. As such, each image represents a time instance along the timeline 520 at which the sensor module 415 retrieves an image from the pixel array 410a.

[0095] In the example frame rate 505, the VBI of the sensor module 415 for the pixel array 410a is set to three time units (e.g., by the controller module 475). That is, the sensor module 415 retrieves an image from the pixel array 410a every third time instance along the timeline 520. As shown in the example frame rate 505, the sensor module 415 retrieves an image at the time instances 525, 540, and 555. Thus, the example frame rate 505 has a frame rate of one image per three time units.

[0096] The example frame rate 510 is similar to the example frame rate 505 except the VBI is set to two time units. Thus, the sensor module 415 retrieves an image from the pixel array 410a every second time instance along the timeline 520. The example frame rate 510 shows the sensor module 415 retrieving an image from the pixel array 410a at the time instances 525, 535, 545, and 555. Since the VBI of the example frame rate 510 is less than the VBI of the example frame rate 505, the frame rate of the example frame rate 510 is higher than the frame rate of the example frame rate 505.

[0097] The example frame rate 515 is also similar to the example frame rate 505 except the VBI of the sensor module 415 for the pixel array 410a is set to one time unit. Therefore, the sensor module 415 is instructed to retrieve an image from the pixel array 410a every time instance along the timeline 520. As illustrated, the sensor module 415 retrieves an image from the pixel array 410a at the time instances 525-555. The VBI of the example frame rate 515 is less than the VBIs of the example frame rates 505 and 510. Therefore, the frame rate of the example frame rate 515 is higher than the example frame rates 505 and 510.

III. Video Conferencing

[0098] A. Video Conference Architecture

[0099] FIG. 6 conceptually illustrates a software architecture for a video conferencing and processing module 600 of a dual camera mobile device of some embodiments. The video conferencing and processing module 600 includes a CIPU driver 605, a media exchange module 610, and an encoder driver 620 that are similar to the corresponding modules and drivers 305, 310, and 320 described above by reference to FIG. 3. The video conferencing and processing module 600 also includes a video conference module 625, a video conference client 645, and a network interface 650 for performing a variety of video conferencing functions. Like the video processing and encoding module 300, the video conferencing and processing module 600 processes and encodes images that are captured from cameras of the dual camera mobile device.

[0100] As described above by reference to FIG. 3, the media exchange module 610 allows consumers and producers of media content in the device to exchange media content and instructions regarding the processing of the media content, the CIPU driver 605 serves as a communication interface with the captured image processing unit (CIPU) 655, and the encoder driver 620 serves as a communication interface with the encoder hardware 660 (e.g., an encoder chip, an encoding component on a system on chip, etc.).

[0101] The video conference module 625 of some embodiments handles various video conferencing functions such as image processing, video conference management, and net-

working. As shown, the video conference module 625 interacts with the media exchange module 610, the video conference client 645, and the network interface 650. In some embodiments, the video conference module 625 receives instructions from and sends instructions to the video conference client 645. The video conference module 625 of some embodiments also sends data to and receives data from networks (e.g., a local area network (LAN), a wireless local area network (WLAN), a wide area network (WAN), a network of networks, a code division multiple access (CDMA) network, a GSM network, etc.) through the network interface 650.

[0102] The video conference module 625 includes an image processing layer 630, a management layer 635, and a network layer 640. In some embodiments, the image processing layer 630 performs image processing operations on images for video conferencing. For example, the image processing layer 630 of some embodiments performs exposure adjustment, image resizing, perspective correction, and dynamic range adjustment as described in further detail below. The image processing layer 630 of some embodiments sends requests through the media exchange module 610 for images from the CIPU 655.

[0103] The management layer 635 of some embodiments controls the operation of the video conference module 625. For instance, in some embodiments, the management layer 635 initializes a camera/cameras of the dual camera mobile device, processes images and audio to transmit to a remote device, and processes images and audio received from the remote device. In some embodiments, the management layer 635 generates composite (e.g., PIP) displays for the device. Moreover, the management layer 635 may change the operation of the video conference module 625 based on networking reports received from the network layer 640.

[0104] In some embodiments, the network layer 640 performs some or all of the networking functionalities for video conferencing. For instance, the network layer 640 of some embodiments establishes a network connection (not shown) between the dual camera mobile device and a remote device of a video conference, transmits images to the remote device, and receives images from the remote device, among other functionalities, as described below and in the above-incorporated U.S. patent application Ser. No. _____, entitled "Establishing Video Conference During a Phone Call," with Attorney Docket Number APLE.P0212. In addition, the network layer 640 receives networking data such as packet loss, one-way latency, and roundtrip delay time, among other types of data, processes such data, and reports the data to the management layer 635.

[0105] The video conference client 645 of some embodiments is an application that may use the video conferencing functions of the video conference module 625 such as a video conferencing application, a voice-over-IP (VoIP) application (e.g., Skype), or an instant messaging application. In some embodiments, the video conference client 645 is a stand-alone application while in other embodiments the video conference client 645 is integrated into another application.

[0106] In some embodiments, the network interface 650 is a communication interface that allows the video conference module 625 and the video conference client 645 to send data and receive data over a network (e.g., a cellular network, a local area network, a wireless network, a network of networks, the Internet, etc.) through the network interface 650. For instance, if the video conference module 625 wants to send data (e.g., images captured by cameras of the dual cam-

era mobile device) to another device on the Internet, the video conference module 625 sends the images to the other device through the network interface 650.

[0107] B. Video Conference Set Up

[0108] FIG. 7 conceptually illustrates an example video conference request messaging sequence 700 of some embodiments. This figure shows the video conference request messaging sequence 700 among a video conference client 710 running on a device 705, a video conference server 715, and a video conference client 725 running on a device 720. In some embodiments, the video conference clients 710 and 725 are the same as the video conference client 645 shown in FIG. 6. As shown in FIG. 7, one device (i.e., the device 705) requests a video conference and another device (i.e., the device 720) responds to such request. The dual camera mobile device described in the present application can perform both operations (i.e., make a request and respond to a request).

[0109] The video conference server 715 of some embodiments routes messages among video conference clients. While some embodiments implement the video conference server 715 on one computing device, other embodiments implement the video conference server 715 on multiple computing devices. In some embodiments, the video conference server is a publicly accessible server that can handle and route messages for numerous conferences at once. Each of the video conference clients 710 and 725 of some embodiments communicates with the video conference server 715 over a network (e.g., a cellular network, a local area network, a wireless network, a network of networks, the Internet etc.) through a network interface such as the network interface 650 described above.

[0110] The video conference request messaging sequence 700 of some embodiments starts when the video conference client 710 receives (at operation 1) a request from a user of the device 705 to start a video conference with the device 720. The video conference client 710 of some embodiments receives the request to start the video conference when the user of the device 705 selects a user interface (UI) item of a user interface displayed on the device 705. Examples of such user interfaces are illustrated in FIG. 8 and FIG. 11, which are described below.

[0111] After the video conference client 710 receives the request, the video conference client 710 sends (at operation 2) a video conference request, which indicates the device 720 as the recipient based on input from the user, to the video conference server 715. The video conference server 715 forwards (at operation 3) the video conference request to the video conference client 725 of the device 720. In some embodiments, the video conference server 715 forwards the video conference request to the video conference client 725 using push technology. That is, the video conference server 715 initiates the transmission of the video conference request to the video conference client 725 upon receipt from the video conference client 710, rather than waiting for the client 725 to send a request for any messages.

[0112] When the video conference client 725 of some embodiments receives the video conference request, a user interface is displayed on the device 720 to indicate to the user of the device 720 that the user of the device 705 sent a request to start a video conference and to prompt the user of the device 720 to accept or reject the video conference request. An example of such a user interface is illustrated in FIG. 9, which is described below. In some embodiments, when the video conference client 725 receives (at operation 4) a request to

accept the video conference request from the user of the device **705**, the video conference client **725** sends (at operation **5**) a video conference acceptance to the video conference server **715**. The video conference client **725** of some embodiments receives the request to accept the video conference request when the user of the device **720** selects a user interface item of a user interface as illustrated in FIG. **9**, for example.

[**0113**] After the video conference server **715** receives the video conference acceptance from the video conference client **725**, the video conference server **715** forwards (at operation **6**) the video conference acceptance to the video conference client **710**. Some embodiments of the video conference server **715** forward the video conference acceptance to the video conference client **710** using the push technology described above.

[**0114**] Upon receiving the video conference acceptance, some embodiments establish (at operation **7**) a video conference between the device **705** and the device **720**. Different embodiments establish the video conference differently. For example, the video conference establishment of some embodiments includes negotiating a connection between the device **705** and the device **720**, determining a bit rate at which to encode video, and exchanging video between the device **705** and the device **720**.

[**0115**] In the above example, the user of the device **720** accepts the video conference request. In some embodiments, the device **720** can be configured (e.g., through the preference settings of the device) to automatically accept incoming video conference requests without displaying a UI. Moreover, the user of the device **720** can also reject (at operation **4**) the video conference request (e.g., by selecting a user interface item of a user interface displayed on the device **720**). Instead of sending a video conference acceptance, the video conference client **725** sends a video conference rejection to the video conference server **715**, which forwards the video conference rejection to the video conference client **710**. The video conference is then never established.

[**0116**] 1. Video Conference Setup User Interface

[**0117**] In some embodiments, a video conference is initiated based on an ongoing phone call. That is, while the user of a mobile device is engaged in a phone call with a second user, the user can turn the phone call into a video conference with the permission of the other party. For some embodiments of the invention, FIG. **8** illustrates the start of such a video conference by a dual camera handheld mobile device **800**. This figure illustrates the start of the video conference in terms of five operational stages **810**, **815**, **820**, **825**, and **830** of a user interface (“UI”) **805** of the device **800**.

[**0118**] As shown in FIG. **8**, the UI **805** includes a name field **835**, a selection menu **840**, and a selectable UI item **845**. The name field **835** displays the name of the person on the other end of the phone call, with whom a user would like to request a video conference. In this example, the selectable UI item **845** (which can be implemented as a selectable button) provides a selectable End Call option for the user to end the phone call. The selection menu **840** displays a menu of selectable UI items, such as a Speakerphone item **842**, a Mute item **844**, a Keypad item **846**, a Phonebook item **848**, a Hold item **852**, a Video Conference item **854**, etc. Different embodiments display the selection menu differently. For the embodiments illustrated by FIG. **8**, the selection menu **840** includes several equally sized icons, each of which represents a differ-

ent operation. Other embodiments provide a scrollable menu, or give priority to particular items (e.g., by making the items larger).

[**0119**] The operation of the UI **805** will now be described by reference to the state of this UI during the five stages, **810**, **815**, **820**, **825**, and **830** that are illustrated in FIG. **8**. In the first stage **810**, a phone call has been established between the handheld mobile device user and Nancy Jones. The second stage **815** displays the UI **805** after the user selects the selectable Video Conference option **854** (e.g., through a single finger tap by finger **850**) to activate a video conference tool. In this example, the Video Conference option **854** (which can be implemented as a selectable icon) allows the user to start a video conference during the phone call. In the second stage, the Video Conference option **854** is highlighted to indicate that the video conference tool has been activated. Different embodiments may indicate such a selection in different ways (e.g., by highlighting the border or the text of the item).

[**0120**] The third stage **820** displays the UI **805** after the device **800** has started the video conference process with the selection of the Video Conference option **854**. The third stage is a transitional hold stage while the device waits for the video conference to be established (e.g., while the device waits for the device on the other end of the call to accept or reject the video conference). In the third stage **820**, the user of the device **800** can still talk to the user of the other device (i.e., Nancy Jones) while the video conference connection is being established. In addition, some embodiments allow the user of the device **800** to cancel the video conference request in the third stage **820** by selecting a selectable UI item displayed on the UI **805** (not shown) for canceling the video conference request. During this hold stage, different embodiments use different displays in the UI **805** to indicate the wait state.

[**0121**] As shown in FIG. **8**, in some embodiments the wait state of the third stage is illustrated in terms of a full screen display of a video being captured by the device **800** along with a “Preview” notation at the bottom of this video. Specifically, in FIG. **8**, the third stage **820** illustrates the start of the video conference process by displaying in a display area **860** of the UI **805** a full screen presentation of the video being captured by the device’s camera. In some embodiments, the front camera is the default camera selected by the device at the start of a video conference. Often, this front camera points to the user of the device at the start of the video conference. Accordingly, in the example illustrated in FIG. **8**, the third stage **820** illustrates the device **800** as presenting a full screen video of the user of the device **800**. The wait state of the device is further highlighted by the “Preview” designation **865** below the video appearing in the display area **860** during the third stage **820**.

[**0122**] The transitional third hold stage **820** can be represented differently in some embodiments. For instance, some embodiments allow the user of the device **800** to select the back camera as the camera for starting the video conference. To allow for this selection, some embodiments allow the user to specify (e.g., through a menu preference setting) the back camera as the default camera for the start of a video conference, and/or allow the user to select the back camera from a menu that displays the back and front cameras after the user selects the Video Conference option **854**. In either of these situations, the UI **805** (e.g., display area **860**) displays a video captured by the back camera during the third hold stage **820**.

[**0123**] Also, other embodiments might indicate the activation of the video conference tool by displaying the smaller

version of the video captured by the device **800**, by displaying a still image that is stored on the device **800**, by providing a message to highlight the wait state of the device (e.g., by showing “Conference Being Established”), by not displaying the “Preview” designation, etc. Also, in the third stage **820**, the UI **805** of some embodiments provides an End button (not shown) to allow the user to cancel entering the video conference and revert back to the phone call if he decides not to enter the video conference at this stage (e.g., while the user is waiting for the remote user to respond to his request).

[0124] The fourth stage **825** illustrates the UI **805** in a transitional state after the remote user has accepted the video conference request and a video conference connection has been established. In this transitional state, the display area **860** that displays the video of the local user (that is being captured by the front camera in this example) gradually decreases in size (i.e., gradually shrinks), as indicated by the arrows **875**. The display area **860** (i.e., the local user’s video) shrinks so that the UI **805** can display a display area **870** (e.g., a display window **870**) that contains the video from a camera of the remote device behind the display area **860**. In other words, the shrinking of the local user’s video **860** creates a PIP display **880** that has a foreground inset display **860** of the local user’s video and a background main display **870** of the remote user. In this example, the background main display **870** presents a video of a lady whose video is being captured by the remote device’s front camera (e.g., Nancy Jones, the user of the remote device) or a lady whose video is being captured by the remote device’s back camera (e.g., a lady whose video is being captured by Nancy Jones). One of ordinary skill will realize that the transitional fourth stage shown in FIG. **8** is simply one exemplary approach used by some embodiments, and that other embodiments might animate the transitional fourth stage differently.

[0125] The fourth stage **825** also illustrates a selectable UI item **832** in a lower display area **855**. The selectable UI item **832** (which can be implemented as a selectable button) provides a selectable End Conference option **832** below the PIP display **880**. The user may select this End Conference option **832** to end the video conference (e.g., through a single finger tap). Different embodiments may allow the user to end the conference in different ways, such as by toggling a switch on the mobile device, by giving voice commands, etc. Moreover, different embodiments may allow the End Conference option **832** to fade away during the video conference, thereby allowing the PIP display **880** to take up the entire display area **885**. The End Conference option **832** may then reappear at a single finger tap at the bottom of the display area **885**, giving the user access to the End Conference option **832**. In some embodiments, the layout of the display area **855** is same as the display area **855** described in further detail below.

[0126] The fifth stage **830** illustrates the UI **805** after the animation of the fourth transitional state **825** has ended. Specifically, the fifth stage **830** illustrates a PIP display **880** that is presented by the UI **805** during the video conference. As mentioned above, this PIP display **880** includes two video displays: a larger background display **870** from the remote camera and a smaller foreground inset display **860** from the local camera.

[0127] This PIP display **880** is only one manner of presenting a composite view of the videos being captured by the remote and local devices. In addition to this composite view, the devices of some embodiments provide other composite views. For example, instead of having a larger background

display **870** of the remote user, the larger background display **870** can be of the local user and the smaller foreground inset display **860** of the remote user. As further described below, some embodiments allow a user to switch during a video conference between the local cameras and/or remote cameras as the cameras for the inset and main views in the PIP display **880**.

[0128] Also, some embodiments allow the local and remote videos to appear in the UI **805** in two side-by-side display areas (e.g., left and right display windows, or top and bottom display windows) or two diagonally aligned display areas. The manner of the PIP display or a default display mode may be specified by the user in some embodiments through the preference settings of the device or through controls that the user can select during a video conference, as further described below and in the above-incorporated U.S. patent application Ser. No. _____, entitled “Establishing Video Conference During a Phone Call,” with Attorney Docket Number APLE. P0212.

[0129] When the user of the device **800** of FIG. **8** invites the remote user to a video conference, the remote user may accept or reject the invitation. FIG. **9** illustrates a UI **905** of the remote user’s device **900** at six different stages **910**, **915**, **920**, **925**, **930**, and **935** that show the sequence of operations for presenting and accepting a video conference invitation at the remote user’s device. The description of the UI **905** below refers to the user of the device **900** (i.e., the device that receives the video conference request) as the invite recipient, and the user of the device **800** (i.e., the device that sends the video conference request) as the invite requestor. Also, in this example, it is assumed that the invite recipient’s device **900** is a dual camera device, like that of the invite requestor. However, in other examples, one or both of these devices are single camera devices.

[0130] The first stage **910** illustrates the UI **905** when the invite recipient receives an invitation to a video conference from the invite requestor, John Smith. As shown in FIG. **9**, the UI **905** in this stage includes a name field **995**, a message field **940**, and two selectable UI items **945** and **950**. The name field **995** displays the name of a person who is requesting a video conference. In some embodiments, the name field **995** displays a phone number of the person who is requesting a video conference instead of the name of the person. The message field **940** displays an invite from the invite requestor to the invite recipient. In this example, the “Video Conference Invitation” in the field **940** indicates that the invite requestor is requesting a video conference with the invite recipient. The selectable UI items **945** and **950** (which can be implemented as selectable buttons) provide selectable Deny Request and Accept Request options **945** and **950** for the invite recipient to use to reject or accept the invitation. Different embodiments may display these options differently and/or display other options.

[0131] Upon seeing the “Video Conference Invitation” notation displayed in the message field **940**, the invite recipient may deny or accept the request by selecting the Deny Request option **945** or Accept Request option **950** in the UI, respectively. The second stage **915** illustrates that in the example shown in FIG. **9**, the user selects the Accept Request option **950**. In this example, this selection is made by the user’s finger tapping on the Accept Request option **950**, and this selection is indicated through the highlighting of this option **950**. Other techniques are provided in some embodiments to select the Accept or Deny Request options **945** and

950 (e.g., double-tapping, etc.) to indicate the selection (e.g., highlighting the border or text of the UI item).

[0132] The third stage 920 displays the UI 905 after the invite recipient has agreed to join the video conference. In this stage, the UI 905 enters into a preview mode that shows a full screen presentation of the video from the remote device's front camera in a display area 944. The front camera in this case is pointed to the user of the remote device (i.e., Nancy Jones in this example). Accordingly, her image is shown in this preview mode. This preview mode allows the invite recipient to make sure that her video is displayed properly and that she is happy with her appearance before the video conference begins (e.g., before actual transmission of the video begins). In some embodiments, a notation, such as a "Preview" notation, may be displayed below the display area 944 to indicate that the invite recipient is in the preview mode.

[0133] Some embodiments allow the invite recipient to select the back camera as the default camera for the start of the video conference, or to select the front or back camera at the beginning of the video conference, as further described in the above-incorporated U.S. patent application Ser. No. _____, entitled "Establishing Video Conference During a Phone Call," with Attorney Docket Number APLE.P0212. Also, other embodiments display the preview display of the invite recipient differently (e.g., in a smaller image placed in the corner of the display area 944). Yet other embodiments do not include this preview mode, but rather start the video conference immediately after the invite recipient accepts the request.

[0134] In the third stage, the UI 905 shows two selectable UI items 975 and 946, one of which overlaps the display area 944 while the other is below this display area 944. The selectable UI item 975 is an Accept button 975 that the user may select to start video conferencing. The selectable UI item 946 is an End button 946 that the invite recipient can select if she decides not to join the video conference at this stage.

[0135] The fourth stage 925 displays the UI 905 after the invite recipient selects the Accept button 975. In this example, the Accept button 975 is highlighted to indicate that the invite recipient is ready to start the video conference. Such a selection may be indicated in different ways in other embodiments.

[0136] The fifth stage 930 illustrates the UI 905 in a transitional state after the invite recipient has accepted the video conference request. In this transitional stage, the display area 944 that displays the video of the invite recipient (that is being captured by the front camera in this example) gradually decreases in size (i.e., gradually shrinks), as indicated by the arrows 960. The invite recipient's video shrinks so that the UI 905 can display a display area 965 (e.g., a display window 965) that contains the video from a camera of the invite requestor behind the display area 944. In other words, the shrinking of the invite recipient's video creates a PIP display 980 that has a foreground inset display area 944 of the invite recipient's video and a background main display 965 of the invite requestor.

[0137] In this example, the background main display 965 presents a video of a man whose video is being captured by the local device's front camera (i.e., John Smith, the user of the local device 800). In another example, this video could have been that of a man whose video is being captured by the local device's back camera (e.g., a man whose video is being captured by John Smith). Different embodiments may animate this transitional fifth stage differently.

[0138] The UI at the fifth stage 930 also displays a display area 855 (e.g., a tool bar or a menu bar) that includes selectable UI item 985 (e.g., mute button 985) for muting the audio of the other user during the video conference, selectable UI item 987 (e.g., end conference button 987) for ending the video conference, and selectable UI item 989 (e.g., switch camera button 989) for switching cameras, which is described in further detail below. As such, the invite recipient may select any of the selectable UI items 985-989 (e.g., through a single finger tap) to perform the desired operation during the video conference. Different embodiments may allow the invite recipient to perform any of the operations in different ways, e.g., by toggling a switch on the mobile device, by giving voice commands, etc.

[0139] Although FIG. 9 shows an example layout for the display area 855, some embodiments provide different layouts of the display area 855 such as the layout of display area 855 of FIG. 8, which includes just a selectable End Conference UI item 832 for ending the video conference. Other layouts of display area 855 can include any number of different selectable UI items for performing different functions. Moreover, the fifth stage 930 shows the display area 855 displayed at the bottom of the UI 905. Different embodiments of the display area 855 can be displayed at different locations within the UI 905 and/or defined as different shapes.

[0140] FIG. 9 shows the display area 855 as a static display area (i.e., the display area 855 is always displayed). However, in some embodiments the display area 855 is a dynamic display area. In some such embodiments, the display area 855 is not ordinarily displayed. Rather, the display area 855 is displayed when a triggering event is received (e.g., a user selection such tapping the display area 980 once, a voice command, etc.). The display area 855 disappears after a user selection is received (e.g., selecting the selectable mute UI item 985) or a defined amount of time (e.g., 3 seconds), which can be specified by the user through the preference settings of the mobile device or the video conference application. In some such embodiments, the display area 855 is automatically displayed after the video conference starts and disappears in the same manner mentioned above.

[0141] The sixth stage 935 illustrates the UI 905 after the animation of the fifth transitional stage has ended. Specifically, the sixth stage illustrates a PIP display 980 that is presented by the UI 905 during the video conference. As mentioned above, this PIP display 980 includes two video displays: a larger background display 965 from the local camera and a smaller foreground inset display 944 from the remote camera. This PIP display 980 is only one manner of presenting a composite view of the videos being captured by the remote and local devices. In addition to this composite view, the devices of some embodiments provide other composite views. For example, instead of having a larger background display of the invite recipient, the larger background display can be of the invite requestor and the smaller foreground inset display of the invite recipient. As further described in the above-incorporated U.S. patent application Ser. No. _____, entitled "Establishing Video Conference During a Phone Call," with Attorney Docket Number APLE.P0212, some embodiments allow a user to control the inset and main views in a PIP display to switchably display the local and remote cameras. Also, some embodiments allow the local and remote videos to appear in the UI 905 in two side-by-side display areas (e.g., left and right display windows, or top and bottom display windows) or two diagonally aligned

display areas. The manner of PIP display or a default display mode may be specified by the user in some embodiments through the preference settings of the device or through controls that the user can select during a video conference, as further described in the above-incorporated U.S. patent application Ser. No. _____, entitled “Establishing Video Conference During a Phone Call,” with Attorney Docket Number APLE.P0212.

[0142] Although FIG. 9 shows the sequence of operations for presenting and accepting a video conference invitation in terms of six different operational stages, some embodiments may implement the operation in less stages. For instance, some of such embodiments may omit presenting the third and fourth stages 920 and 925 and go from the second stage 915 to the fifth stage 930 after the user selects the Accept Request option 950. Other embodiments that implement that operation (i.e., presenting and accepting a video conference invitation) in less stages may omit the first and second stages 910 and 915 and present the user with the third stage 920 when the invite recipient receives an invitation to a video conference from the invite requestor.

[0143] FIG. 10 illustrates an example of performing the operation illustrated in FIG. 9 in less stages by combining the first and third stages into one stage and the second and fourth stage into one stage. In particular, this figure illustrates a UI 905 of the remote user's device 900 at five different stages 1090, 1092, 1094, 930, and 935. The first stage 1090 is similar to the stage 810 except the name field 995 displays the name “John Smith” to indicate the name of the person on the other end of the telephone call. That is, a phone call has been established between the user of the remote mobile device and the user of the local device (i.e., John Smith in this example). The second and third stages 1092 and 1094 are similar to the first and second stages 910 and 915 of FIG. 9 except the second and third stage 1092 and 1094 also show a preview of the user of the remote mobile device (i.e., Nancy Jones in this example). The fourth and fifth stages 930 and 935 are the same as the fifth and sixth stages 930 and 935 of FIG. 9.

[0144] In addition to activating the video conference tool through a selectable option during a phone call, some embodiments allow a user of a dual camera device to initiate a video conference directly without having to make a phone call first. FIG. 11 illustrates another such alternative method to initiate a video conference. This figure illustrates the UI 1105 at seven different stages 1110, 1115, 1120, 1125, 1130, 1135, and 1140 that show an alternative sequence of operations for starting a video conference.

[0145] In the first stage 1110, a user is looking through a contacts list on this mobile device for the person with whom he wants to engage in a video conference, similar to how he would find a contact to call. In the second stage 1115, the user selects the person 1155 with whom he would like to have a video conference (e.g., through a single finger tap 1160 on the person's name 1155). This selection triggers the UI 1105 to display the contact's information and various user selectable options. In this example, Jason's name 1155 is highlighted to indicate that this is the person with whom the user would like to have a video conference. Different embodiments may indicate such a selection in different ways. While the second stage 1115 allows the user of the device 1100 to select a person with whom the user would like to have a video conference through a contact list, some embodiments allow the user to select the person through a “Recents” call history that lists a particular

number or name of a person with whom the user of the device 1100 recently had a video conference or a phone call.

[0146] In the third stage 1120, the UI 1105 displays the selected person's information 1162 and various selectable UI items 1168, 1172, and 1170 after the person's name 1155 has been selected. In this example, one of the various selectable UI items 1172 (which can be implemented as a selectable icon or button) provides a video conference tool. The Video Conference option 1172 allows the user to invite the person identified by the contact 1166 to a video conference. Different embodiments display the information 1162 and selectable UI items 1168, 1172, and 1170 differently (e.g., in a different arrangement).

[0147] The fourth stage 1125 shows the user selecting the Video Conference option 1172 (e.g., through a single finger tap). In this example, the Video Conference option 1172 is highlighted to indicate that the video conference tool 1172 has been activated. Such selections may be indicated differently in different embodiments (e.g., by highlighting the text or border of the selected icon).

[0148] The fifth, sixth and seventh stages 1130, 1135, and 1140 are similar to the third, fourth and fifth stages 820, 825, and 830 illustrated in FIG. 8 and may be understood by reference to the discussion of those stages. In brief, the fifth stage 1130 illustrates a transitional holding stage that waits for the remote user to respond to the invitation to a video conference. The sixth stage 1135 illustrates that after the remote user has accepted the video conference request, the display area 1180 (that displays the video of the local user) gradually decreases in size so the UI 1105 can show a display area 1192 that contains the video from a camera of the remote user behind the display area 1180. In the seventh stage 1140, the PIP display 1147 is presented by the UI 1105 during the video conference. In some embodiments, the layout of display area 855 in the sixth stage 1135 and the seventh stage 1140 is like the layout of the display area 855 of FIG. 9, described above.

[0149] FIGS. 7, 8, 9, 10, and 11 show several ways of establishing a video conference. In some embodiments, during a telephone call, audio data (e.g., voice) is transmitted through one communication channel (over a communication network like a circuit-switched communication network or a packet-switched communication network) and, during a video conference, audio data is transmitted through another communication channel. Thus, in such embodiments, audio data (e.g., voice) is transmitted through a communication channel before the video conference is established, and once the video conference is established, audio is transmitted through a different communication channel (instead of the communication channel used during the telephone call).

[0150] In order to provide a seamless transition (e.g., hand-off) of audio data from the telephone call to the video conference, some embodiments do not terminate the telephone call before establishing the video conference. For instance, some embodiments establish a peer-to-peer video conference connection (e.g., after completing the message sequence illustrated in FIG. 7) before terminating the phone call and starting to transmit audio/video data through the peer-to-peer communication session. Alternatively, other embodiments establish a peer-to-peer video conference connection (e.g., after completing the message sequence illustrated in FIG. 7) and start transmitting audio/video data through the peer-to-peer communication session, before terminating the phone call and starting to present the received audio/video data.

[0151] A peer-to-peer video conference connection of some embodiments allows the mobile devices in the video conference to directly communicate with each other (instead of communicating through a central server, for example). Some embodiments of a peer-to-peer video conference allow the mobile devices in the video conferences to share resources with each other. For instance, through a control communication channel of a video conference, one mobile device can remotely control operations of another mobile device in the video conference by sending instructions from the one mobile device to the other mobile device to direct the other mobile device to process images differently (i.e., share its image processing resource) such as an exposure adjustment operation, a focus adjustment operation, and/or a switch camera operation, described in further detail below.

[0152] 2. Dynamic Bit Rate Setup

[0153] Typically, mobile devices in a video conference communicate data (e.g., audio and video images) to each other through communication channels of different types of communication networks such as different private and public wireless communication networks (e.g., cellular networks like GSM, UMTS, etc.). Examples of such wireless communication networks will be explained by reference to FIGS. 21 and 22, below.

[0154] A communication network's available bandwidth for conducting a video conference differ at different times due to the varying number of mobile devices accessing the communication network at a given time. The available bandwidth can even change during the video conference. In addition, flooding the communication network with high bit rates during the video conference or large amounts of signaling in an attempt to figure out an optimal video conference bit rate is undesirable.

[0155] Because of these reasons, some embodiments employ a novel method for specifying the initial optimal bit rate for the video conference. To identify the initial optimal bit rate for the video conference, this method starts a video conference at a particular bit rate and incrementally increases the bit rate at particular intervals if these embodiments do not detect network conditions that would degrade the video conference quality.

[0156] One example of such embodiments is illustrated in FIG. 12. This figure conceptually illustrates a process 1200 of some embodiments of the invention for setting the bit rate for a video conference. The process 1200 is performed as part of a video conference setup (e.g., as part of the video conference establishment illustrated in FIG. 7) to dynamically determine a bit rate for transmitting data (e.g., audio and video images) based on various network conditions. In some embodiments, the process 1200 is performed by the management layer 635 of the video conference module 625 described above by reference to FIG. 6. A more detailed version of this video conference module will be described below by reference to FIG. 13.

[0157] As shown in FIG. 12, the process 1200 starts by setting (at 1205) a bit rate at an initial bit rate. In some embodiments, the initial bit rate is a default baseline rate for the device. Some embodiments, though, allow a user to specify the initial bit rate. At 1205, the process 1200 also starts the video conference transmitting data (e.g., audio and video images) through one or more communication channels at the initial bit rate to a remote device.

[0158] Next, the process 1200 identifies (at 1210) a set of network condition parameters received from the remote

device in the video conference. In some embodiments, the local device receives the set of network condition parameters from the remote device through a real-time transport protocol (RTP) communication session established at the start of the video conference. For instance, some embodiments provide the network condition parameters through the extension feature of RTP. Moreover, the RTP extension feature of some embodiments can be used to transmit any type of information (such as the set of network condition parameters) by indicating the presence of an extension header in an RTP packet header and defining an extension header for the additional information.

[0159] In different embodiments, devices in the video conference relay different sets of network condition/congestion parameters. In the embodiments described below, the set of network condition parameters include one-way latency and a bandwidth estimation bit rate. In other embodiments, the set of network condition parameters includes packet loss data and roundtrip time (RTT) delay data. As such, different embodiments can include any number of different network condition parameters in the set of network condition parameters.

[0160] In some embodiments, the set of network condition parameters received from the remote device of the video conference is based on data (e.g., audio and video) transmitted from the local mobile device (i.e., the mobile device performing the process 1200) to the remote device during the video conference at the initial bit rate set at operation 1205. For example, in some embodiments, the remote device can determine one-way latency by calculating the time it takes for audio packets to propagate through the network connection from the local mobile device to the remote device by using timestamps of the audio packets. Specifically, in some embodiments, each audio packet is time stamped. In the absence of packet delay, the remote devices should receive the audio packets at set intervals that are equal to the difference in the time stamps. However, when there is one-way latency delay, the remote device receives the audio packets at intervals greater than the difference in the time stamps.

[0161] Also, in some embodiments, the remote device determines a bandwidth estimation bit rate by examining the time a video packet is received, the time a consecutive video packet is received, and the size of the consecutive video packet. That is, the difference in time between the receipt of the two consecutive video packets and the size of the second video packet is used to estimate the available bandwidth of the network connection. Some embodiments determine a bandwidth estimation bit rate by examining multiple pairs of consecutive video packets. The above examples use specific types of data (i.e., audio data for determining one-way latency and video data for determining a bandwidth estimation bit rate). However, other types of data communicated through the network connection between the local mobile device and the remote device can be used as well in some embodiments.

[0162] After identifying (at 1210) the set of network conditions, the process 1200 then determines (at 1215) whether the one-way latency has deteriorated past a defined threshold amount. In some embodiments, the threshold amount is defined as a particular latency amount and the one-way latency is determined to deteriorate past the threshold amount when the difference between the current one-way latency and a previous one-way latency exceeds the particular latency amount. In other embodiments, the threshold amount is defined as a particular rate of change of one-way latencies. As

such, the one-way latency is determined to deteriorate past the threshold amount when the rate of change of a set of one-way latencies (e.g., the current one-way latency and previous one-way latencies) exceeds the particular rate of change.

[0163] When the one-way latency is determined to deteriorate past the threshold amount, the process **1200** ends. Otherwise, the process **1200** determines (at **1220**) whether the current bit rate has reached the bandwidth estimation bit rate. In some embodiments, the bandwidth estimation bit rate indicates an amount of available bandwidth (e.g., 15 kilobits/sec (kbps)) for the network connection. When the process **1200** determines that the current bit rate exceeds the bandwidth estimation bit rate, the process **1200** ends. When the process **1200** determines that the current bit rate does not exceed the bandwidth estimation bit rate, the process **1200** proceeds to operation **1225**.

[0164] At **1225**, the process **1200** determines whether the current bit rate has reached a defined maximum bit rate. When the process **1200** determines that the current bit rate exceeds the defined maximum bit rate, the process **1200** ends. Otherwise, the process **1200** proceeds to operation **1230** to increase the current bit rate by a defined amount. Different embodiments define the amount to increase the bit rate differently. Examples of a defined amount to increase the current bit rate include 32 kbps, 64 kbps, among any number of other amounts to increase the bit rate.

[0165] Next, the process determines (at **1235**) whether a defined amount of time has elapsed. The defined amount of time can be 1 second, 2 seconds, 5 seconds, or any other possible amounts of time since different embodiments define the amount of time differently. The process **1200** waits for the defined amount of time to elapse in order for the remote device to receive data (e.g., audio and video images) transmitted from the local mobile device at the newly increased bit rate (at operation **1230**) and determine network condition parameters based on the newly increased bit rate. If the process **1200** determines that the defined amount of time has not elapsed, the process **1200** returns to operation **1235** until the defined amount of time has elapsed. When the process **1200** determines that the defined amount of time has elapsed, the process **1200** returns to operation **1210**. The operation of the process **1200** from **1210** onwards proceeds as described above until the process **1200** ends.

[0166] When the process **1200** ends (i.e., after operation **1215**, **1220**, or **1225**), the setup of a bit rate for the video conference is complete and an optimal bit rate has been determined. Since the available bandwidth for the video conference can change during the video conference, some embodiments continue to adjust the bit rate based on the set of network condition parameters (i.e., the one-way latency and the bandwidth estimation bit rate) that are received from the remote device. The bit rate can be adjusted during the video conference by increasing the bit rate. For instance, if the process **1200** ended because the one-way latency deteriorated past the defined threshold amount and, during the video conference, the one-way latency improves, some embodiments increase the bit rate. Similarly, if the process **1200** ended because the bit rate exceeded the bandwidth estimation bit rate and, during the video conference, the bandwidth estimation bit rate increases, some embodiments increase the bit rate.

[0167] In contrast, the bit rate can be adjusted during the video conference by decreasing the bit rate. For example, if

the one-way latency continues to deteriorate past the defined threshold amount during the video conference, some embodiments decrease the bit rate. Also, if the bit rate continues to exceed the bandwidth estimation bit rate (e.g., the bandwidth estimation bit rate continues to decrease) during the video conference, some embodiments decrease the bit rate.

[0168] Moreover, the description of the process **1200** uses one-way latency and a bandwidth estimation bit rate to determine whether to increase the bit rate. However, one of ordinary skill will realize that any number of network condition parameters can be used to determine whether to increase the bit rate in different embodiments. For instance, determining whether to increase the bit rate can be based on solely RTT delay data or packet loss data.

[0169] C. Video Conference Architecture

[0170] As mentioned above, FIG. **13** conceptually illustrates a software architecture for a video conferencing and processing module **1300** of a dual camera mobile device of some embodiments. As shown, the video conferencing and processing module **1300** includes a client application **1365**, a video conference module **1302**, a media exchange module **1320**, a buffer **1325**, a captured image processing unit (CIPU) driver **1330**, an encoder driver **1335**, and a decoder driver **1340**. In some embodiments, the buffer **1325** is a frame buffer that stores images of a video for display on a display **1345** of the dual camera mobile device.

[0171] In some embodiments, the client application **1365** is the same as the video conference client **645** of FIG. **6**. As mentioned above, the client application **1365** may be integrated into another application or implemented as a stand-alone application. The client application **1365** may be an application that uses the video conferencing functions of the video conference module **1302**, such as a video conferencing application, a voice-over-IP (VoIP) application (e.g., Skype), or an instant messaging application.

[0172] The client application **1365** of some embodiments sends instructions to the video conference module **1302** such as instructions to start a conference and end a conference, receives instructions from the video conference module **1302**, routes instructions from a user of the dual camera mobile device to the video conference module **1302**, and generates user interfaces that are displayed on the dual camera mobile device and allow a user to interact with the application.

[0173] D. Video Conference Manager

[0174] As shown in FIG. **13**, the video conference module **1302** includes a video conference manager **1304**, an image processing manager **1308**, a networking manager **1314**, and buffers **1306**, **1310**, **1312**, **1316**, and **1318**. In some embodiments, the video conference module **1302** is the same as the video conference module **625** illustrated in FIG. **6** and thus performs some or all of the same functions described above for the video conference module **625**.

[0175] In some embodiments, the video conference manager **1304** is responsible for initializing some or all of the other modules of the video conference module **1302** (e.g., the image processing manager **1308** and the networking manager **1314**) when a video conference is starting, controlling the operation of the video conference module **1302** during the video conference, and ceasing the operation of some or all of the other modules of the video conference module **1302** when the video conference is ending.

[0176] The video conference manager **1304** of some embodiments also processes images received from one or more devices in the video conference and images captured by

one of both cameras of the dual camera mobile device for display on the dual camera mobile device. For instance, the video conference manager 1304 of some embodiments retrieves decoded images, that were received from another device participating in the video conference, from the buffer 1318 and retrieves images processed by CIPU 1350 (i.e., images captured by the dual camera mobile device) from the buffer 1306. In some embodiments, the video conference manager 1304 also scales and composites the images before displaying the images on the dual camera mobile device. That is, the video conference manager 1304 generates the PIP or other composite views to display on the mobile device in some embodiments. Some embodiments scale the images retrieved from the buffers 1306 and 1318 while other embodiments just scale images retrieved from one of the buffers 1306 and 1318.

[0177] Although FIG. 13 illustrates the video conference manager 1304 as part of the video conference module 1302, some embodiments of the video conference manager 1304 are implemented as a component separate from the video conference module 1302. As such, a single video conference manager 1304 can be used to manage and control several video conference modules 1302. For instance, some embodiments will run a separate video conference module on the local device to interact with each party in a multi-party conference, and each of these video conference modules on the local device are managed and controlled by the one video conference manager.

[0178] The image processing manager 1308 of some embodiments processes images captured by the cameras of the dual camera mobile device before the images are encoded by the encoder 1355. For example, some embodiments of the image processing manager 1308 perform one or more of exposure adjustment, focus adjustment, perspective correction, dynamic range adjustment, and image resizing on images processed by the CIPU 1350. In some embodiments, the image processing manager 1308 controls the frame rate of encoded images that are transmitted to the other device in the video conference.

[0179] Some embodiments of the networking manager 1314 manage one or more connections between the dual camera mobile device and the other device participating in the video conference. For example, the networking manager 1314 of some embodiments establishes the connections between the dual camera mobile device and the other device of the video conference at the start of the video conference and tears down these connections at the end of the video conference.

[0180] During the video conference, the networking manager 1314 transmits images encoded by the encoder 1355 to the other device of the video conference and routes images received from the other device of the video conference to decoder 1360 for decoding. In some embodiments, the networking manager 1314, rather than the image processing manager 1308, controls the frame rate of the images that are transmitted to the other device of the video conference. For example, some such embodiments of the networking manager 1314 control the frame rate by dropping (i.e., not transmitting) some of the encoded frames that are supposed to be transmitted to the other device of the video conference.

[0181] As shown, the media exchange module 1320 of some embodiments includes a camera source module 1322, a video compressor module 1324, and a video decompressor module 1326. The media exchange module 1320 is the same

as the media exchange module 310 shown in FIG. 3, with more detail provided. The camera source module 1322 routes messages and media content between the video conference module 1302 and the CIPU 1350 through the CIPU driver 1330, the video compressor module 1324 routes message and media content between the video conference module 1302 and the encoder 1355 through the encoder driver 1335, and the video decompressor module 1326 routes messages and media content between the video conference module 1302 and the decoder 1360 through the decoder driver 1340. Some embodiments implement the TNR module 315 included in the media exchange module 310 (not shown in FIG. 13) as part of the camera source module 1322 while other embodiments implement the TNR module 315 as part of the video compressor module 1324.

[0182] In some embodiments, the CIPU driver 1330 and the encoder driver 1335 are the same as the CIPU driver 305 and the encoder driver 320 illustrated in FIG. 3. The decoder driver 1340 of some embodiments acts as a communication interface between the video decompressor module 1326 and decoder 1360. In such embodiments, the decoder 1360 decodes images received from the other device of the video conference through the networking manager 1314 and routed through the video decompressor module 1326. After the images are decoded, they are sent back to the video conference module 1302 through the decoder driver 1340 and the video decompressor module 1326.

[0183] In addition to performing video processing during a video conference, the video conferencing and processing module 1300 for the dual camera mobile device of some embodiments also performs audio processing operations during the video conference. FIG. 14 illustrates such a software architecture. As shown, the video conferencing and processing module 1300 includes the video conference module 1302 (which includes the video conference manager 1304, the image processing manager 1308, and the networking manager 1314), the media exchange module 1320, and the client application 1365. Other components and modules of the video conferencing and processing module 1300 shown in FIG. 13 are omitted in FIG. 14 to simplify the description. The video conferencing and processing module 1300 also includes frame buffers 1405 and 1410, audio processing manager 1415, and audio driver 1420. In some embodiments, the audio processing manager 1415 is implemented as a separate software module while in other embodiments the audio processing manager 1415 is implemented as part of the media exchange module 1320.

[0184] The audio processing manager 1415 processes audio data captured by the dual camera mobile device for transmission to the other device in the video conference. For example, the audio processing manager 1415 receives audio data through the audio driver 1420, which is captured by microphone 1425, and encodes the audio data before storing the encoded audio data in the buffer 1405 for transmission to the other device. The audio processing manager 1415 also processes audio data captured by and received from the other device in the video conference. For instance, the audio processing manager 1415 retrieves audio data from the buffer 1410 and decodes the audio data, which is then output through the audio driver 1420 to the speaker 1430.

[0185] In some embodiments, the video conference module 1302 along with the audio processing manager 1415 and its associated buffers are part of a larger conference module. When a multi-participant audio conference is conducted

between several devices without exchange of video content, this video conferencing and processing module **1300** only uses the networking manager **1314** and the audio processing manager **1415** to facilitate the exchange of audio over an Internet Protocol (IP) layer.

[0186] The operation of the video conference manager **1304** of some embodiments will now be described by reference to FIG. 15. FIG. 15 conceptually illustrates a process **1500** performed by a video conference manager of some embodiments such as video conference manager **1304** illustrated in FIG. 13. This can be equivalent to being performed by the management layer **635** of FIG. 6. In some embodiments, the video conference manager **1304** performs process **1500** when a user of the dual camera mobile device accepts (e.g., through a user interface displayed on the dual camera mobile device) a video conference request or when a user of another device accepts a request sent by the user of the dual camera mobile device.

[0187] The process **1500** begins by receiving (at **1505**) instructions to start a video conference. In some embodiments, the instructions are received from the client application **1365** or are received from a user through a user interface displayed on the dual camera mobile device and forwarded to the video conference manager **1304** by the client application **1365**. For example, in some embodiments, when a user of the dual camera mobile device accepts a video conference request, the instructions are received through the user interface and forwarded by the client application. On the other hand, when a user of the other device accepts a request sent from the local device, some embodiments receive the instructions from the client application without user interface interaction (although there may have been previous user interface interaction to send out the initial request).

[0188] Next, the process **1500** initializes (at **1510**) a first module that interacts with the video conference manager **1304**. The modules of some embodiments that interact with the video conference manager **1304** include the CIPU **1350**, the image processing manager **1308**, the audio processing manager **1415**, and the networking manager **1314**.

[0189] In some embodiments, initializing the CIPU **1350** includes instructing the CIPU **1350** to start processing images captured by one or both cameras of the dual camera mobile device. Some embodiments initialize the image processing manager **1308** by instructing the image processing manager **1308** to start retrieving images from the buffer **1310** and processing and encoding the retrieved images. To initialize the audio processing manager **1415**, some embodiments instruct the audio processing manager **1415** to begin encoding audio data captured by the microphone **1425** and decoding audio data stored in the buffer **1410** (which was received from the other device) in order to output to the speaker **1430**. The initializing of the networking manager **1314** of some embodiments includes instructing the networking manager **1314** to establish a network connection with the other device in the video conference.

[0190] The process **1500** then determines (at **1515**) whether there are any modules left to initialize. When there are modules left to initialize, the process **1500** returns to operation **1510** to initialize another of the modules. When all of the required modules have been initialized, the process **1500** generates (at **1520**) composite images for displaying on the dual camera mobile device (i.e., local display). These composite images may include those illustrated in FIG. 65 in the above-incorporated U.S. patent application Ser. No. _____,

entitled "Establishing Video Conference During a Phone Call," with Attorney Docket Number APLE.P0212, and can include various combinations of images from the cameras of the local dual camera mobile device and images from cameras of the other device participating in the video conference.

[0191] Next, the process **1500** determines (at **1525**) whether a change has been made to the video conference. Some embodiments receive changes to the video conference through user interactions with a user interface displayed on the dual camera mobile device while other embodiments receive changes to the video conference from the other device through the networking manager **1314** (i.e., remote control). The changes to video conference settings may also be received from the client application **1365** or other modules in the video conference module **1302** in some embodiments. The video conference settings may also change due to changes in the network conditions.

[0192] When a change has been made, the process **1500** determines (at **1530**) whether the change to the video conference is a change to a network setting. In some embodiments, the changes are either network setting changes or image capture setting changes. When the change to the video conference is a change to a network setting, the process modifies (at **1540**) the network setting and then proceeds to operation **1545**. Network setting changes of some embodiments include changing the bit rate at which images are encoded or the frame rate at which the images are transmitted to the other device.

[0193] When the change to the video conference is not a change to a network setting, the process **1500** determines that the change is a change to an image capture setting and then proceeds to operation **1535**. The process **1500** then performs (at **1535**) the change to the image capture setting. In some embodiments, change to the image capture settings may include switching cameras (i.e., switching which camera on the dual camera mobile device will capture video), focus adjustment, exposure adjustment, displaying or not displaying images from one or both cameras of the dual camera mobile device, and zooming in or out of images displayed on the dual camera mobile device, among other setting changes.

[0194] At operation **1545**, the process **1500** determines whether to end the video conference. When the process **1500** determines to not end the video conference, the process **1500** returns to operation **1520**. When the process **1500** determines that the video conference will end, the process **1500** ends. Some embodiments of the process **1500** determine to end the video conference when the process **1500** receives instructions from the client application **1365** to end the video conference (i.e., due to instructions received through the user interface of the local dual camera mobile device or received from the other device participating in the video conference).

[0195] In some embodiments, the video conference manager **1304** performs various operations when the video conference ends that are not shown in process **1500**. Some embodiments instruct the CIPU **1350** to stop producing images, the networking manager **1314** to tear down the network connection with the other device in the video conference, and the image processing manager **1308** to stop processing and encoding images.

[0196] E. Image Processing Manager & Encoder

[0197] In addition to temporal noise reduction and image processing operations performed by the CIPU and/or CIPU driver, some embodiments perform a variety of image processing operations at the image processing layer **630** of the

video conference module **625**. These image processing operations may include exposure adjustment, focus adjustment, perspective correction, adjustment of dynamic range, and image resizing, among others.

[0198] FIG. 16 conceptually illustrates a process **1600** for performing such image processing operations. In some embodiments, some or all of the operations of the process **1600** are performed by a combination of the image processing manager **1308** and the encoder driver **1335** of FIG. 13. In some of such embodiments, the image processing manager **1308** performs the pixel-based processing (e.g., resizing, dynamic range adjustment, perspective correction, etc.). Some embodiments perform process **1600** during a video conference on images that are to be transmitted to another device participating in the video conference.

[0199] The process **1600** will now be described by reference to FIG. 13. The process starts by retrieving (at **1605**) an image from the buffer **1306**. In some embodiments, the retrieved image is an image of a video (i.e., an image in a sequence of images). This video may have been captured by a camera of a device on which the process **1600** is performed.

[0200] Next, the process **1600** performs (at **1610**) exposure adjustment on the retrieved image. Some embodiments perform exposure adjustments through a user interface that is displayed on the dual camera mobile device. The process **1600** next performs (at **1615**) focus adjustment on the image. Some embodiments perform focus adjustment through a user interface that is displayed on the dual camera mobile device.

[0201] The process **1600** performs (at **1620**) image resizing on the image. Some embodiments perform image resizing on the image to reduce the number of bits used to encode the image (i.e., lower the bit rate). In some embodiments, the process **1600** performs image resizing as described by reference to FIG. 26 in the above-incorporated U.S. patent application Ser. No. _____, entitled "Establishing Video Conference During a Phone Call," with Attorney Docket Number APLE.P0212.

[0202] The process **1600** next performs (at **1625**) perspective correction on the image. In some embodiments, the process **1600** performs perspective correction as described in FIG. 24 in the above-incorporated U.S. patent application Ser. No. _____, entitled "Establishing Video Conference During a Phone Call," with Attorney Docket Number APLE.P0212. Such perspective correction involves using data taken by one or more accelerometer and/or gyroscope sensors that identifies orientation and movement of the dual camera mobile device. This data is then used to modify the image to correct for the perspective being off.

[0203] After perspective correction is performed on the image, the process **1600** adjusts (at **1630**) the dynamic range of the image. In some embodiments, the dynamic range of an image is the range of possible values that each pixel in the image can have. For example, an image with a dynamic range of 0-255 can be adjusted to a range of 0-128 or any other range of values. Adjusting the dynamic range of an image can reduce the amount of bits that will be used to encode the image (i.e., lower the bit rate) and thereby smooth out the image.

[0204] Adjusting the dynamic range of an image can also be used for various other purposes. One purpose is to reduce image noise (e.g., the image was captured by a noisy camera sensor). To reduce noise, the dynamic range of the image can be adjusted so that the black levels are redefined to include lighter blacks (i.e., crush blacks). In this manner, the noise of

the image is reduced. Another purpose of dynamic range adjustment is to adjust one or more colors or range of colors in order to enhance the image. For instance, some embodiments may assume that the image captured by the front camera is an image of a person's face. Accordingly, the dynamic range of the image can be adjusted to increase the red and pink colors to make the person's cheeks appear rosy/rosier. The dynamic range adjustment operation can be used for other purposes as well.

[0205] Finally, the process **1600** determines (at **1635**) one or more rate controller parameters that are used to encode the image. Such rate controller parameters may include a quantization parameter and a frame type (e.g., predictive, bi-directional, intra-coded) in some embodiments. The process then ends.

[0206] While the various operations of process **1600** are illustrated as being performed in a specific order, one of ordinary skill will recognize that many of these operations (exposure adjustment, focus adjustment, perspective correction, etc.) can be performed in any order and are not dependent on one another. That is, the process of some embodiments could perform focus adjustment before exposure adjustment, or similar modifications to the process illustrated in FIG. 16.

[0207] F. Networking Manager

[0208] FIG. 17 conceptually illustrates the software architecture of a networking manager **1700** of some embodiments such as the networking manager **1314** illustrated in FIG. 13. As described above, the networking manager **1700** manages network connections (e.g., connection establishment, connection monitoring, connection adjustments, connection tear down, etc.) between a dual camera mobile device on which it operates and a remote device in a video conference. During the video conference, the networking manager **1700** of some embodiments also processes data for transmission to the remote device and processes data received from the remote device.

[0209] As shown in FIG. 17, the networking manager **1700** includes a session negotiating manager **1705**, a transmitter module **1715**, a universal transmission buffer **1720**, a universal transmission buffer manager **1722**, a virtual transport protocol (VTP) manager **1725**, a receiver module **1730**, and a media transport manager **1735**.

[0210] The session negotiating manager **1705** includes a protocol manager **1710**. The protocol manager **1710** ensures that the transmitter module **1715** uses a correct communication protocol to transmit data to a remote device during the video conference and enforces rules of the communication protocol that is used. Some embodiments of the protocol manager **1710** support a number of communication protocols, such as a real-time transport protocol (RTP), a transmission control protocol (TCP), a user datagram protocol (UDP), and a hypertext transfer protocol (HTTP), among others.

[0211] The session negotiating manager **1705** is responsible for establishing connections between the dual camera mobile device and one or more remote devices participating in the video conference, as well as tearing down these connections after the conference. In some embodiments, the session negotiating manager **1705** is also responsible for establishing multimedia communication sessions (e.g., to transmit and receive video and/or audio streams) between the dual camera mobile device and the remote devices in the video conference (e.g., using a session initiation protocol (SIP)).

[0212] The session negotiating manager 1705 also receives feedback data from the media transport manager 1735 and, based on the feedback data, determines the operation of the universal transmission buffer 1720 (e.g., whether to transmit or drop packets/frames) through the universal transmission buffer manager 1722. This feedback, in some embodiments, may include one-way latency and a bandwidth estimation bit rate. In other embodiments, the feedback includes packet loss information and roundtrip delay time (e.g., determined based on packets sent to the remote device in the video conference and the receipt of acknowledgements from that device). Based on the information from the media transport manager 1735, the session negotiating manager 1705 can determine whether too many packets are being sent and instruct the universal transmission buffer manager 1722 to have the universal transmission buffer 1720 transmit fewer packets (i.e., to adjust the bit rate as described in FIG. 12).

[0213] The transmitter module 1715 retrieves encoded images (e.g., as a bitstream) from a video buffer (e.g., the buffer 1312 of FIG. 13) and packetizes the images for transmission to a remote device in the video conference through the universal transmission buffer 1720 and the virtual transport protocol manager 1725. The manner in which the encoded images are created and sent to the transmitter module 1715 can be based on instructions or data received from the media transport manager 1735 and/or the session negotiating manager 1705. In some embodiments, packetizing the images involves breaking the received bitstream into a group of packets each having a particular size (i.e., a size specified by the session negotiating manager 1705 according to a particular protocol), and adding any required headers (e.g., address headers, protocol specification headers, etc.).

[0214] The universal transmission buffer manager 1722 controls the operation of the universal transmission buffer 1720 based on data and/or instructions received from the session negotiating manager 1705. For example, the universal transmission buffer manager 1722 may be instructed to direct the universal transmission buffer 1720 to transmit data, stop transmitting data, drop data, etc. As described above, in some embodiments when a remote device participating in the conference appears to be dropping packets, this will be recognized based on acknowledgements received from the remote device. To reduce the packet dropping, the universal transmission buffer manager 1722 may be instructed to transmit packets at a slower rate to the remote device.

[0215] The universal transmission buffer 1720 stores data received from the transmitter module 1715 and transmits the data to the remote device through the VTP manager 1725. As noted above, the universal transmission buffer 1720 may drop data (e.g., images of the video) based on instructions received from the universal transmission buffer manager 1722.

[0216] In some embodiments, RTP is used to communicate data packets (e.g., audio packets and video packets) over UDP during a video conference. Other embodiments use RTP to communicate data packets over TCP during the video conference. Other transport layer protocols can be used as well in different embodiments.

[0217] Some embodiments define a particular communication channel between two mobile devices by a pair of port numbers (i.e., source port number and destination port number). For instance, one communication channel between the mobile devices can be defined by one pair of port numbers (e.g., source port 50 and destination port 100) and another different communication channel between the mobile devices

can be defined by another different pair of port numbers (e.g., source port 75 and destination port 150). Some embodiments also use a pair of internet protocol (IP) addresses in defining communication channels. Different communication channels are used to transmit different types of data packets in some embodiments. For example, video data packets, audio data packets, and control signaling data packets can be transmitted in separate communication channels. As such, a video communication channel transports video data packets and an audio communication channel transports audio data packets.

[0218] In some embodiments, a control communication channel is for messaging between the local mobile device and a remote device during a video conference. Examples of such messaging include sending and receiving requests, notifications, and acknowledgements to such requests and notifications. Another example of messaging includes sending remote control instruction messages from one device to another. For instance, the remote control operations described in the above-incorporated U.S. patent application Ser. No. _____, entitled "Establishing Video Conference During a Phone Call," with Attorney Docket Number APLE.P0212 (e.g., instructing a device to only send images from one particular camera or to only capture images with a particular camera) can be performed by sending instructions from a local device to a remote device through the control communication channel for the local device to remotely control operations of the remote device. Different embodiments implement the control communication using different protocols like a real-time transport control protocol (RTCP), an RTP extension, SIP, etc. For instance, some embodiments use RTP extension to relay one set of control messages between two mobile devices in a video conference and use SIP packets to relay another set of control messages between the mobile devices during the video conference.

[0219] The VTP manager 1725 of some embodiments allows different types of data packets that are specified to be transmitted through different communication channels (e.g., using different pairs of port numbers) to be transmitted through a single communication channel (e.g., using the same pair of port numbers). One technique for doing this involves identifying the data packet types, identifying the communication channel through which data packets are specified to be transmitted by extracting the specified pair of port numbers of the data packets, and specifying the data packets to be transmitted through the single communication channel by modifying the pair of port numbers of the data packets to be the pair of port numbers of the single communication channel (i.e., all the data packets are transmitted through the same pair of port numbers).

[0220] To keep track of the original pair of port numbers for each type of data packet, some embodiments store a mapping of the original pair of port numbers for the data packet type. Some of these embodiments than use the packet type field of the protocol to differentiate the different packets that are being multiplexed into one communication channel. For instance, some embodiments that have the VTP manager multiplex audio, video and control packets into one RTP stream, use the RTP packet type field to differentiate between the audio, video and control packets that are transmitted in the one RTP channel to the other device in the video conference. In some of these embodiments, the VTP manager also routes control messaging in SIP packets to the other device.

[0221] Some embodiments identify examine the data packet signatures (i.e., packet header formats) to distinguish

between different packets that are communicated using different protocols (e.g., to differentiate between packets transported using RTP and packets transported using SIP). In such embodiments, after the data packets of the different protocols are determined, the fields of the data packets that use the same protocol (e.g., audio data and video data using RTP) are examined as described above to identify the different data types. In this manner, the VTP manager 1725 transmits different data packets, which are intended to be transmitted through different communication channels, through a single communication channel.

[0222] Although one way of combining different types of data through a single communication channel is described above, other embodiments utilize other techniques to multiplex different packet types into one communication stream. For example, one technique of some embodiments involves keeping track of the original pair of port numbers of the data packets and storing the original pair of port numbers in the data packet itself to be later extracted. Still other ways exist for combining different types of data between two video conference participants into one port pair channel.

[0223] When the VTP manager 1725 receives data packets from the remote device through a virtualized communication channel, the VTP manager 1725 examines the signatures of the data packets to identify the different packets that are sent using the different protocols. Such signatures can be used to differentiate SIP packets from RTP packets. The VTP manager of some embodiments also uses the packet type field of some or all of the packets to demultiplex the various different types of packets (e.g., audio, video and control packets) that were multiplexed into a single virtualized channel. After identifying these different types of packets, the VTP manager associates each different type of packet with its corresponding port pair numbers based on a mapping of port pair numbers and packet types that it keeps. The VTP manager 1725 then modifies the pair of port numbers of the data packets with the identified pair of port numbers and forwards the data packets to be depacketized. In other embodiments that use different techniques for multiplexing different packet types into the single channel, the VTP manager uses different techniques for parsing out the packets.

[0224] By using such techniques for multiplexing and demultiplexing the different packets, the VTP manager 1725 creates a single virtualized communication channel (e.g., a single pair of port numbers), transmits the video data, audio data, and control signaling data through the single virtualized communication channel, and receives audio, video, and control packets from the remote device through the single virtualized communication channel. Thus, from the perspective of the network, data is transmitted through this single virtualized communication channel, while, from the perspective of the session negotiating manager 1705 and the protocol manager 1710, the video data, audio data, and control signaling data are transmitted through different communication channels.

[0225] Similar to the images that are transmitted to the remote device in the video conference, images transmitted from the remote device in the video conference are received in packet format. The receiver module 1730 receives the packets and depacketizes them in order to reconstruct the images before storing the images in a video buffer (e.g., the buffer 1316 of FIG. 13) to be decoded. In some embodiments, depacketizing the images involves removing any headers and reconstructing a bitstream that only has image data (and potentially size data) from the packets.

[0226] The media transport manager 1735 processes feedback data (e.g., one-way latency, bandwidth estimation bit rate, packet loss data, roundtrip delay time data, etc.) received from the network to dynamically and adaptively adjust the rate of data transmission (i.e., bit rate). The media transport manager 1735 also controls error resilience based on the processed feedback data in some other embodiments, and may also send the feedback data to the video conference manager 1304 in order to adjust other operations of the video conference module 1302 such as scaling, resizing, and encoding. In addition to having the universal transmission buffer drop packets when a remote device in the conference is not able to process all of the packets, the video conference module and encoder can use a lower bit rate for encoding the images so that fewer packets will be sent for each image.

[0227] In some embodiments, the media transport manager 1735 may also monitor other variables of the device such as power consumption and thermal levels that may affect how the operational power modes of the cameras are configured, as discussed above. This data may also be used as additional inputs into the feedback data (e.g., if the device is getting too hot, the media transport manager 1735 may try to have the processing slowed down).

[0228] Several example operations of the networking manager 1700 will now be described by reference to FIG. 13. The transmission of images captured by a camera of the dual camera mobile device to a remote device in the video conference will be described first, followed by the description of receiving images from the remote device. The transmitter module 1715 retrieves encoded images from the buffer 1312, which are to be transmitted to the remote device in the video conference.

[0229] The protocol manager 1710 determines the appropriate protocol to use (e.g., RTP to transmit audio and video) and the session negotiating manager 1705 informs the transmitter module 1715 of such protocol. Next, the transmitter module 1715 packetizes the images and sends the packetized images to the universal transmission buffer 1720. The universal transmission buffer manager 1722 receives instructions from the session negotiating manager 1705 to direct the universal transmission buffer 1720 to transmit or drop the images. The VTP manager 1725 receives the packets from the universal transmission buffer 1720 and processes the packets in order to transmit the packets through a single communication channel to the remote device.

[0230] When receiving images from the remote device, the VTP manager 1725 receives packetized images from the remote device through the virtualized single communication channel and processes the packets in order to direct the images to the receiver module 1730 through a communication channel that is assigned to receive the images (e.g., a video communication channel).

[0231] The receiver module 1730 depacketizes the packets to reconstruct the images and sends the images to the buffer 1316 for decoding by the decoder 1360. The receiver module 1730 also forwards control signaling messages to the media transport manager 1735 (e.g., acknowledgements of received packets from the remote device in the video conference).

[0232] Several example operations of the networking manager 1700 were described above. These are only illustrative examples, as various other embodiments will perform these or different operations using different modules or with functionalities spread differently between the modules. Furthermore, additional operations such as dynamic bit rate adjust-

ment may be performed by the modules of networking manager 1700 or other modules.

IV. Electronic System

[0233] Many of the above-described features and applications are implemented as software processes that are specified as a set of instructions recorded on a computer readable storage medium (also referred to as computer readable medium). When these instructions are executed by one or more processing unit(s) (e.g., one or more processors, cores of processors, or other processing units), they cause the processing unit(s) to perform the actions indicated in the instructions. Examples of computer readable media include, but are not limited to, CD-ROMs, flash drives, RAM chips, hard drives, EPROMs, etc. The computer readable media does not include carrier waves and electronic signals passing wirelessly or over wired connections.

[0234] In this specification, the term “software” is meant to include firmware residing in read-only memory or applications stored in magnetic storage which can be read into memory for processing by a processor. Also, in some embodiments, multiple software inventions can be implemented as sub-parts of a larger program while remaining distinct software inventions. In some embodiments, multiple software inventions can also be implemented as separate programs. Finally, any combination of separate programs that together implement a software invention described here is within the scope of the invention. In some embodiments, the software programs, when installed to operate on one or more electronic systems, define one or more specific machine implementations that execute and perform the operations of the software programs.

[0235] Some embodiments are implemented as software processes that include one or more application programming interfaces (APIs) in an environment with calling program code interacting with other program code being called through the one or more interfaces. Various function calls, messages or other types of invocations, which further may include various kinds of parameters, can be transferred via the APIs between the calling program and the code being called. In addition, an API may provide the calling program code the ability to use data types or classes defined in the API and implemented in the called program code.

[0236] At least certain embodiments include an environment with a calling software component interacting with a called software component through an API. A method for operating through an API in this environment includes transferring one or more function calls, messages, other types of invocations or parameters via the API.

[0237] One or more Application Programming Interfaces (APIs) may be used in some embodiments. For example, some embodiments of the media exchange module 310 (or 910) provide a set of APIs to other software components for accessing various video processing and encoding functionalities described in FIGS. 3 and 6.

[0238] An API is an interface implemented by a program code component or hardware component (hereinafter “API-implementing component”) that allows a different program code component or hardware component (hereinafter “API-calling component”) to access and use one or more functions, methods, procedures, data structures, classes, and/or other services provided by the API-implementing component. An

API can define one or more parameters that are passed between the API-calling component and the API-implementing component.

[0239] An API allows a developer of an API-calling component (which may be a third party developer) to leverage specified features provided by an API-implementing component. There may be one API-calling component or there may be more than one such component. An API can be a source code interface that a computer system or program library provides in order to support requests for services from an application. An operating system (OS) can have multiple APIs to allow applications running on the OS to call one or more of those APIs, and a service (such as a program library) can have multiple APIs to allow an application that uses the service to call one or more of those APIs. An API can be specified in terms of a programming language that can be interpreted or compiled when an application is built.

[0240] In some embodiments the API-implementing component may provide more than one API, each providing a different view of or with different aspects that access different aspects of the functionality implemented by the API-implementing component. For example, one API of an API-implementing component can provide a first set of functions and can be exposed to third party developers, and another API of the API-implementing component can be hidden (not exposed) and provide a subset of the first set of functions and also provide another set of functions, such as testing or debugging functions which are not in the first set of functions. In other embodiments the API-implementing component may itself call one or more other components via an underlying API and thus be both an API-calling component and an API-implementing component.

[0241] An API defines the language and parameters that API-calling components use when accessing and using specified features of the API-implementing component. For example, an API-calling component accesses the specified features of the API-implementing component through one or more API calls or invocations (embodied for example by function or method calls) exposed by the API and passes data and control information using parameters via the API calls or invocations. The API-implementing component may return a value through the API in response to an API call from an API-calling component. While the API defines the syntax and result of an API call (e.g., how to invoke the API call and what the API call does), the API may not reveal how the API call accomplishes the function specified by the API call. Various API calls are transferred via the one or more application programming interfaces between the calling (API-calling component) and an API-implementing component. Transferring the API calls may include issuing, initiating, invoking, calling, receiving, returning, or responding to the function calls or messages; in other words, transferring can describe actions by either of the API-calling component or the API-implementing component. The function calls or other invocations of the API may send or receive one or more parameters through a parameter list or other structure. A parameter can be a constant, key, data structure, object, object class, variable, data type, pointer, array, list or a pointer to a function or method or another way to reference a data or other item to be passed via the API.

[0242] Furthermore, data types or classes may be provided by the API and implemented by the API-implementing component. Thus, the API-calling component may declare vari-

ables, use pointers to, use or instantiate constant values of such types or classes by using definitions provided in the API.

[0243] Generally, an API can be used to access a service or data provided by the API-implementing component or to initiate performance of an operation or computation provided by the API-implementing component. By way of example, the API-implementing component and the API-calling component may each be any one of an operating system, a library, a device driver, an API, an application program, or other module (it should be understood that the API-implementing component and the API-calling component may be the same or different type of module from each other). API-implementing components may in some cases be embodied at least in part in firmware, microcode, or other hardware logic. In some embodiments, an API may allow a client program to use the services provided by a Software Development Kit (SDK) library. In other embodiments an application or other client program may use an API provided by an Application Framework. In these embodiments the application or client program may incorporate calls to functions or methods provided by the SDK and provided by the API or use data types or objects defined in the SDK and provided by the API. An Application Framework may in these embodiments provide a main event loop for a program that responds to various events defined by the Framework. The API allows the application to specify the events and the responses to the events using the Application Framework. In some implementations, an API call can report to an application the capabilities or state of a hardware device, including those related to aspects such as input capabilities and state, output capabilities and state, processing capability, power state, storage capacity and state, communications capability, etc., and the API may be implemented in part by firmware, microcode, or other low level logic that executes in part on the hardware component.

[0244] The API-calling component may be a local component (i.e., on the same data processing system as the API-implementing component) or a remote component (i.e., on a different data processing system from the API-implementing component) that communicates with the API-implementing component through the API over a network. It should be understood that an API-implementing component may also act as an API-calling component (i.e., it may make API calls to an API exposed by a different API-implementing component) and an API-calling component may also act as an API-implementing component by implementing an API that is exposed to a different API-calling component.

[0245] The API may allow multiple API-calling components written in different programming languages to communicate with the API-implementing component (thus the API may include features for translating calls and returns between the API-implementing component and the API-calling component); however the API may be implemented in terms of a specific programming language. An API-calling component can, in one embodiment, call APIs from different providers such as a set of APIs from an OS provider and another set of APIs from a plug-in provider and another set of APIs from another provider (e.g. the provider of a software library) or creator of the another set of APIs.

[0246] FIG. 18 is a block diagram illustrating an exemplary API architecture, which may be used in some embodiments of the invention. As shown in FIG. 18, the API architecture 1800 includes the API-implementing component 1810 (e.g., an operating system, a library, a device driver, an API, an application program, software or other module) that implements

the API 1820. The API 1820 specifies one or more functions, methods, classes, objects, protocols, data structures, formats and/or other features of the API-implementing component that may be used by the API-calling component 1830. The API 1820 can specify at least one calling convention that specifies how a function in the API-implementing component 1810 receives parameters from the API-calling component 1830 and how the function returns a result to the API-calling component. The API-calling component 1830 (e.g., an operating system, a library, a device driver, an API, an application program, software or other module), makes API calls through the API 1820 to access and use the features of the API-implementing component 1810 that are specified by the API 1820. The API-implementing component 1810 may return a value through the API 1820 to the API-calling component 1830 in response to an API call.

[0247] It will be appreciated that the API-implementing component 1810 may include additional functions, methods, classes, data structures, and/or other features that are not specified through the API 1820 and are not available to the API-calling component 1830. It should be understood that the API-calling component 1830 may be on the same system as the API-implementing component 1810 or may be located remotely and accesses the API-implementing component 1810 using the API 1820 over a network. While FIG. 18 illustrates a single API-calling component 1830 interacting with the API 1820, it should be understood that other API-calling components, which may be written in different languages (or the same language) than the API-calling component 1830, may use the API 1820.

[0248] The API-implementing component 1810, the API 1820, and the API-calling component 1830 may be stored in a machine-readable medium, which includes any mechanism for storing information in a form readable by a machine (e.g., a computer or other data processing system). For example, a machine-readable medium includes magnetic disks, optical disks, random access memory; read only memory, flash memory devices, etc.

[0249] FIG. 19 is an example of a dual camera mobile computing device architecture 1900. The implementation of a mobile computing device can include one or more processing units 1905, memory interface 1910 and a peripherals interface 1915. Each of these components that make up the computing device architecture can be separate components or integrated in one or more integrated circuits. These various components can also be coupled together by one or more communication buses or signal lines.

[0250] The peripherals interface 1915 can be coupled to various sensors and subsystems, including a camera subsystem 1920, a wireless communication subsystem(s) 1925, audio subsystem 1930, I/O subsystem 1935, etc. The peripherals interface 1915 enables communication between processors and peripherals. Peripherals such as an orientation sensor 1945 or an acceleration sensor 1950 can be coupled to the peripherals interface 1915 to facilitate the orientation and acceleration functions.

[0251] The camera subsystem 1920 can be coupled to one or more optical sensors 1940, e.g., a charged coupled device (CCD) optical sensor, a complementary metal-oxide-semiconductor (CMOS) optical sensor. The camera subsystem 1920 coupled with the sensors may facilitate camera functions, such as image and/or video data capturing. Wireless communication subsystems 1925 may serve to facilitate communication functions. Wireless communication subsystems

1925 may include radio frequency receivers and transmitters, and optical receivers and transmitters. They may be implemented to operate over one or more communication networks such as a GSM network, a Wi-Fi network, Bluetooth network, etc. The audio subsystems **1930** is coupled to a speaker and a microphone to facilitate voice-enabled functions, such as voice recognition, digital recording, etc.

[0252] I/O subsystem **1935** involves the transfer between input/output peripheral devices, such as a display, a touch screen, etc., and the data bus of the CPU through the Peripherals Interface. I/O subsystem **1935** can include a touch-screen controller **1955** and other input controllers **1960** to facilitate these functions. Touch-screen controller **1955** can be coupled to the touch screen **1965** and detect contact and movement on the screen using any of multiple touch sensitivity technologies. Other input controllers **1960** can be coupled to other input/control devices, such as one or more buttons.

[0253] Memory interface **1910** can be coupled to memory **1970**, which can include high-speed random access memory and/or non-volatile memory such as flash memory. Memory can store an operating system (OS) **1972**. The OS **1972** can include instructions for handling basic system services and for performing hardware dependent tasks.

[0254] Memory can also include communication instructions **1974** to facilitate communicating with one or more additional devices; graphical user interface instructions **1976** to facilitate graphic user interface processing; image/video processing instructions **1978** to facilitate image/video-related processing and functions; phone instructions **1980** to facilitate phone-related processes and functions; media exchange and processing instructions **1982** to facilitate media communication and processing-related processes and functions; camera instructions **1984** to facilitate camera-related processes and functions; and video conferencing instructions **1986** to facilitate video conferencing processes and functions. The above identified instructions need not be implemented as separate software programs or modules. Various functions of mobile computing device can be implemented in hardware and/or in software, including in one or more signal processing and/or application specific integrated circuits.

[0255] The above-described embodiments may include touch I/O device **2001** that can receive touch input for interacting with computing system **2003**, as shown in FIG. 20, via wired or wireless communication channel **2002**. Touch I/O device **2001** may be used to provide user input to computing system **2003** in lieu of or in combination with other input devices such as a keyboard, mouse, etc. One or more touch I/O devices **2001** may be used for providing user input to computing system **2003**. Touch I/O device **2001** may be an integral part of computing system **2003** (e.g., touch screen on a laptop) or may be separate from computing system **2003**.

[0256] Touch I/O device **2001** may include a touch sensitive panel which is wholly or partially transparent, semitransparent, non-transparent, opaque or any combination thereof. Touch I/O device **2001** may be embodied as a touch screen, touch pad, a touch screen functioning as a touch pad (e.g., a touch screen replacing the touchpad of a laptop), a touch screen or touchpad combined or incorporated with any other input device (e.g., a touch screen or touchpad disposed on a keyboard) or any multi-dimensional object having a touch sensitive surface for receiving touch input.

[0257] In one example, touch I/O device **2001** embodied as a touch screen may include a transparent and/or semitrans-

parent touch sensitive panel partially or wholly positioned over at least a portion of a display. According to this embodiment, touch I/O device **2001** functions to display graphical data transmitted from computing system **2003** (and/or another source) and also functions to receive user input. In other embodiments, touch I/O device **2001** may be embodied as an integrated touch screen where touch sensitive components/devices are integral with display components/devices. In still other embodiments a touch screen may be used as a supplemental or additional display screen for displaying supplemental or the same graphical data as a primary display and receiving touch input.

[0258] Touch I/O device **2001** may be configured to detect the location of one or more touches or near touches on device **2001** based on capacitive, resistive, optical, acoustic, inductive, mechanical, chemical measurements, or any phenomena that can be measured with respect to the occurrences of the one or more touches or near touches in proximity to device **2001**. Software, hardware, firmware or any combination thereof may be used to process the measurements of the detected touches to identify and track one or more gestures. A gesture may correspond to stationary or non-stationary, single or multiple, touches or near touches on touch I/O device **2001**. A gesture may be performed by moving one or more fingers or other objects in a particular manner on touch I/O device **2001** such as tapping, pressing, rocking, scrubbing, twisting, changing orientation, pressing with varying pressure and the like at essentially the same time, contiguously, or consecutively. A gesture may be characterized by, but is not limited to a pinching, sliding, swiping, rotating, flexing, dragging, or tapping motion between or with any other finger or fingers. A single gesture may be performed with one or more hands, by one or more users, or any combination thereof.

[0259] Computing system **2003** may drive a display with graphical data to display a graphical user interface (GUI). The GUI may be configured to receive touch input via touch I/O device **2001**. Embodied as a touch screen, touch I/O device **2001** may display the GUI. Alternatively, the GUI may be displayed on a display separate from touch I/O device **2001**. The GUI may include graphical elements displayed at particular locations within the interface. Graphical elements may include but are not limited to a variety of displayed virtual input devices including virtual scroll wheels, a virtual keyboard, virtual knobs, virtual buttons, any virtual UI, and the like. A user may perform gestures at one or more particular locations on touch I/O device **2001** which may be associated with the graphical elements of the GUI. In other embodiments, the user may perform gestures at one or more locations that are independent of the locations of graphical elements of the GUI. Gestures performed on touch I/O device **2001** may directly or indirectly manipulate, control, modify, move, actuate, initiate or generally affect graphical elements such as cursors, icons, media files, lists, text, all or portions of images, or the like within the GUI. For instance, in the case of a touch screen, a user may directly interact with a graphical element by performing a gesture over the graphical element on the touch screen. Alternatively, a touch pad generally provides indirect interaction. Gestures may also affect non-displayed GUI elements (e.g., causing user interfaces to appear) or may affect other actions within computing system **2003** (e.g., affect a state or mode of a GUI, application, or operating system). Gestures may or may not be performed on touch I/O device **2001** in conjunction with a displayed cursor. For

instance, in the case in which gestures are performed on a touchpad, a cursor (or pointer) may be displayed on a display screen or touch screen and the cursor may be controlled via touch input on the touchpad to interact with graphical objects on the display screen. In other embodiments in which gestures are performed directly on a touch screen, a user may interact directly with objects on the touch screen, with or without a cursor or pointer being displayed on the touch screen.

[0260] Feedback may be provided to the user via communication channel **2002** in response to or based on the touch or near touches on touch I/O device **2001**. Feedback may be transmitted optically, mechanically, electrically, olfactory, acoustically, or the like or any combination thereof and in a variable or non-variable manner.

[0261] These functions described above can be implemented in digital electronic circuitry, in computer software, firmware or hardware. The techniques can be implemented using one or more computer program products. Programmable processors and computers can be included in or packaged as mobile devices. The processes and logic flows may be performed by one or more programmable processors and by one or more programmable logic circuitry. General and special purpose computing devices and storage devices can be interconnected through communication networks.

[0262] Some embodiments include electronic components, such as microprocessors, storage and memory that store computer program instructions in a machine-readable or computer-readable medium (alternatively referred to as computer-readable storage media, machine-readable media, or machine-readable storage media). Some examples of such computer-readable media include RAM, ROM, read-only compact discs (CD-ROM), recordable compact discs (CD-R), rewritable compact discs (CD-RW), read-only digital versatile discs (e.g., DVD-ROM, dual-layer DVD-ROM), a variety of recordable/rewritable DVDs (e.g., DVD-RAM, DVD-RW, DVD+RW, etc.), flash memory (e.g., SD cards, mini-SD cards, micro-SD cards, etc.), magnetic and/or solid state hard drives, read-only and recordable Blu-Ray® discs, ultra density optical discs, any other optical or magnetic media, and floppy disks. The computer-readable media may store a computer program that is executable by at least one processing unit and includes sets of instructions for performing various operations. Examples of computer programs or computer code include machine code, such as is produced by a compiler, and files including higher-level code that are executed by a computer, an electronic component, or a microprocessor using an interpreter.

[0263] While the above discussion primarily refers to microprocessor or multi-core processors that execute software, some embodiments are performed by one or more integrated circuits, such as application specific integrated circuits (ASICs) or field programmable gate arrays (FPGAs). In some embodiments, such integrated circuits execute instructions that are stored on the circuit itself.

[0264] As used in this specification and any claims of this application, the terms “computer”, “server”, “processor”, and “memory” all refer to electronic or other technological devices. These terms exclude people or groups of people. For the purposes of the specification, the terms display or displaying means displaying on an electronic device. As used in this specification and any claims of this application, the terms “computer readable medium” and “computer readable media” are entirely restricted to tangible, physical objects that

store information in a form that is readable by a computer. These terms exclude any wireless signals, wired download signals, and any other ephemeral signals.

[0265] FIG. 21 conceptually illustrates an example communication system **2100** used for connecting some participants of a video conference according to some embodiments. As shown, the communication system **2100** includes several mobile devices **2115**, several cellular base stations (or Node Bs) **2110**, several radio network controllers (RNCs) **2105**, and a core network **2125**. Cellular base stations and RNCs are collectively referred to as a Universal Mobile Telecommunications System (UMTS) Terrestrial Radio Access Network (UTRAN) **2130**. Each RNC **2105** is connected to one or more cellular base stations **2110** that, together, are referred to as a radio access network (RAN).

[0266] Each cellular base station **2110** covers a service region **2120**. As shown, the mobile devices **2115** in each service region are wirelessly connected to the serving cellular base station **2110** of the service region **2120** through a Uu interface. The Uu interface uses a protocol stack that has two planes: a control plane and a user plane. The user plane supports circuit-switched, packet-switched and broadcast data streams. The control plane carries the network's signaling messages.

[0267] Each cellular base station is connected to an RNC through an Iub interface. Each RNC **2105** is connected to the core network **2125** by Iu-cs and an Iu-ps interfaces. The Iu-cs interface is used for circuit switched services (e.g., voice) while the Iu-ps interface is used for packet switched services (e.g., data). The Iur interface is used for connecting two RNCs together.

[0268] Accordingly, the communication system **2100** supports both circuit-switched services and packet-switched services. For example, circuit-switched services allow a telephone call to be conducted by transmitting the telephone call data (e.g., voice) through circuit-switched equipment of the communication system **2100**. Packet-switched services allow a video conference to be conducted by using a transport protocol layer such as UDP or TCP over an internet layer protocol like IP to transmit video conference data through packet-switched equipment of the communication system **2100**. In some embodiments, the telephone call to video conference transition (e.g., handoff) previously described in the Video Conference Setup section uses the circuit-switched and packet-switched services supported by a communication system like the communication system **2100**. That is, in such embodiments, the telephone call is conducted through the circuit-switched equipment of the communication system **2100** and the video conference it conducted through the packet-switched equipment of the communication system **2100**.

[0269] Although the example communication system in FIG. 21 illustrates a third generation (3G) technology UTRAN wireless mobile communication system, it should be noted that second generation (2G) communication systems, other 3 G communication systems such as 3GPP2 Evolution-Data Optimized or Evolution-Data only (EV-DO) and 3rd generation partnership project 2 (3GPP2) Code Division Multiple Access 1× (CDMA 1×), fourth generation (4G) communication systems, wireless local area network (WLAN), and Worldwide Interoperability for Microwave Access (WiMAX) communication systems can be used for connecting some of the participants of a conference in some embodiments. Examples of 2G systems include Global Sys-

tem for Mobile communications (GSM), General Packet Radio Service (GPRS), and Enhanced Data Rates for GSM Evolution (EDGE). A 2 G communication system architecture is similar to the architecture shown in FIG. 21 except the 2 G communication system architecture uses base transceiver stations (BTSs) instead of Node Bs 2110 and base station controllers (BSC) instead of RNC 2105. In a 2 G communication system, an A interface between the BSC and the core network is used for circuit switched services and a Gb interface between the BSC and the core network is used for packet switched services.

[0270] In some embodiments, the communication system 2100 is operated by a service carrier who initially provisions a mobile device 2115 to allow the mobile device 2115 to use the communication system 2100. Some embodiments provision a mobile device 2115 by configuring and registering a subscriber identity module (SIM) card in the mobile device 2115. In other embodiments, the mobile device 2115 is instead configured and registered using the mobile device 2115's memory. Moreover, additional services can be provisioned (after a customer purchases the mobile device 2115) such as data services like GPRS, multimedia messaging service (MMS), and instant messaging. Once provisioned, the mobile device 2115 is activated and is thereby allowed to use the communication system 2100 by the service carrier.

[0271] The communication system 2100 is a private communication network in some embodiments. In such embodiments, the mobile devices 2115 can communicate (e.g., conduct voice calls, exchange data) among each other (e.g., mobile devices 2115 that are provisioned for the communication system 2100). In other embodiments, the communication system 2100 is a public communication network. Thus, the mobile devices 2115 can communicate with other devices outside of the communication system 2100 in addition to the mobile devices 2115 provisioned for the communication system 2100. Some of the other devices outside of the communication system 2100 include phones, computers, and other devices that connect to the communication system 2100 through other networks such as a public switched telephone network or another wireless communication network.

[0272] The Long-Term Evolution (LTE) specification is used to define 4 G communication systems. FIG. 22 conceptually illustrates an example of a 4 G communication system 2200 that is used for connecting some participants of a video conference in some embodiments. As shown, the communication system 2200 includes several mobile devices 2115, several Evolved Node Bs (eNBs) 2205, a Mobility Management Entity (MME) 2215, a Serving Gateway (S-GW) 2220, a Packet Data Network (PDN) Gateway 2225, and a Home Subscriber Server (HSS) 2235. In some embodiments, the communication system 2200 includes one or more MMEs 2215, one or more S-GWs 2220, one or more PDN Gateways 2225, and one or more HSSs 2235.

[0273] The eNBs 2205 provide an air interface for the mobile devices 2115. As shown, each eNB 2205 covers a service region 2210. The mobile devices 2115 in each service region 2210 are wirelessly connected to the eNB 2205 of the service region 2210 through a LTE-Uu interface. FIG. 22 also shows the eNBs 2205 connected to each other through an X2 interface. In addition, the eNBs 2205 are connected to the MME 2215 through an S1-MME interface and to the S-GW 2220 through an S1-U interface. The eNBs 2205 are collectively referred to as an Evolved UTRAN (E-UTRAN) 2230.

[0274] The eNBs 2205 provide functions such as radio resource management (e.g., radio bearer control, connection mobility control, etc.), routing of user plane data towards the S-GW 2220, signal measurement and measurement reporting, MME selection at the time of mobile device attachment, etc. The MME 2215 functions include idle mode mobile device tracking and paging, activation and deactivation of radio bearers, selection of the S-GW 2220 at the time of mobile device attachment, Non-Access Stratum (NAS) signaling termination, user authentication by interacting with the HSS 2235, etc.

[0275] The S-GW 2220 functions includes (1) routing and forwarding user data packets and (2) managing and storing mobile device contexts such as parameters of the IP bearer service and network internal routing information. The PDN Gateway 2225 functions include providing connectivity from the mobile devices to external packet data networks (not shown) by being the point of exit and entry of traffic for the mobile devices. A mobile station may have simultaneous connectivity with more than one PDN Gateway for accessing multiple packet data networks. The PDN Gateway 2225 also acts as the anchor for mobility between 3GPP and non-3GPP technologies such as WiMAX and 3GPP2 (e.g., CDMA 1x and EV-DO).

[0276] As shown, MME 2215 is connected to S-GW 2220 through an S11 interface and to the HSS 2235 through an S6a interface. The S-GW 2220 and the PDN Gateway 2220 are connected through an S8 interface. The MME 2215, S-GW 2220, and PDN Gateway 2225 are collectively referred to as an Evolved Packet Core (EPC). The EPC is the main component of a System Architecture Evolution (SAE) architecture, which is the core network architecture of 3GPP LTE wireless communication standard. The EPC is a pure packet system. For example, the EPC does not have a voice media gateway. Services, like voice and SMS, are packet-switched routed and are provided by application functions that make use of the EPC service. So using the telephone call to video conference transition previously described above as an example, both the telephone call and the video conference are conducted through packet-switched equipment of the communication system 2200 in some embodiments. In some such embodiments, the packet-switched channel used for the telephone call is continued to be used for the audio data of the video conference after the telephone call terminates. However, in other such embodiments, a different packet-switched channel is created (e.g., when the video conference is established) and audio data is transmitted through the newly created packet-switched channel instead of the packet-switched channel of the telephone call when the telephone call terminates.

[0277] Moreover, the amount of bandwidth provided by these different technologies ranges from 44 kilobits per second (kbps) for GPRS to over 10 megabits per second (Mbps) for LTE. Download rates of 100 Mbps and upload rates of 22 Mbps are predicted in the future for LTE.

[0278] While the invention has been described with reference to numerous specific details, one of ordinary skill in the art will recognize that the invention can be embodied in other specific forms without departing from the spirit of the invention. In addition, a number of the figures conceptually illustrate processes. The specific operations of these processes may not be performed in the exact order shown and described. The specific operations may not be performed in one continuous series of operations, and different specific operations may be performed in different embodiments. Furthermore, the

process could be implemented using several sub-processes, or as part of a larger macro process.

[0279] Also, many embodiments were described above by reference to a video conference between two dual camera mobile devices. However, one of ordinary skill in the art will realize that many of these embodiments are used in cases involving a video conference between a dual camera mobile device and another device, such as a single camera mobile device, a computer, a phone with video conference capability, etc. Moreover, many of the embodiments described above can be used in single camera mobile devices and other computing devices with video conference capabilities. Thus, one of ordinary skill in the art would understand that the invention is not limited by the foregoing illustrative details, but rather is to be defined by the appended claims.

We claim:

1. A computer readable medium of a first mobile device, the computer readable medium storing a computer program for managing a video conference between the first device and a second device, the computer program comprising sets of instructions for:

- identifying a first ceiling bit rate for transmitting video conference data to the second device through the communication channel;
- identifying a current bit rate that is less than the first ceiling bit rate;
- receiving networking data regarding the communication channel from the second device;
- from the received network data, determining that the communication channel will sustain an increase in the current bit rate;
- incrementing the current bit rate; and
- iteratively performing said receiving, determining and incrementing operations until a determination is made that the communication channel will not sustain an increase in the current bit rate.

2. The computer readable medium of claim 1, wherein the second device is a mobile device communicating with the first mobile device through a wireless communication network with limited bandwidth.

3. The computer readable medium of claim 1, wherein the sets of instructions for iteratively performing the receiving, determining and incrementing operations is repeatedly executed at a particular temporal rate to allow for accurate network data to be relayed from the second device based on a prior incremental increase of the current bit rate.

4. The computer readable medium of claim 1, wherein the computer program further comprises sets of instructions for even after the iterative performance of the receiving, determining and incrementing operations has terminated, receiving networking data regarding the communication channel from the second device, and from the received network data, identifying a change in the available resources of the communication channel;

incrementing or decrementing the current bit rate based on the identified change in the available resources of the communication channel.

5. The computer readable medium of claim 1, wherein the set of instructions for determining comprises a set of instructions for determining that the current bit rate has reached the first ceiling bit rate.

6. The computer readable medium of claim 1, wherein the received data comprises an available bandwidth measurement data that is provided by the second device.

wherein the set of instructions for determining comprises a set of instructions for determining that the current bit rate has reached a second ceiling bit rate that is lower than the first ceiling bit rate and that is derived from the available bandwidth measurement data provided by the second device.

7. The computer readable medium of claim 1, wherein the received data comprises a one way signal delay that is computed by the second device.

8. The computer readable medium of claim 7, wherein the one way signal delay is computed by the second device by examining delay of audio packets sent from the first device to the second device.

9. The computer readable medium of claim 8, wherein the audio packet delay is examined by examining the difference in time codes in the audio packets with the difference in the reception of the audio packets by the second device.

10. The computer readable medium of claim 1, wherein the received data comprises an available bandwidth measurement data that is computed by the second device by examining the size of video packets transmitted to the second device with the difference in time that the second device receives the packets.

11. The computer readable medium of claim 1, wherein the received network data comprises packet loss data.

12. The computer readable medium of claim 1, wherein the received network data comprises roundtrip delay time data.

13. The computer readable medium of claim 1, wherein the set of instructions for receiving network data comprises a set of instructions for receiving network data in packets that are sent along a real-time transport protocol (RTP) stream and that are identified in the RTP stream by RTP Header Extensions.

14. A method performed by a first device to manage a video conference between the first device and a second device, the method comprising:

- identifying a first ceiling bit rate for transmitting video conference data to the second device through the communication channel;
- identifying a current bit rate that is less than the first ceiling bit rate;
- receiving networking data regarding the communication channel from the second device;
- from the received network data, determining that the communication channel will sustain an increase in the current bit rate;
- incrementing the current bit rate; and
- iteratively performing said receiving, determining and incrementing operations until a determination is made that the communication channel will not sustain an increase in the current bit rate.

15. The method of claim 14, wherein the second device is a mobile device communicating with the first mobile device through a wireless communication network with limited bandwidth.

16. The method of claim 14, wherein iteratively performing the receiving, determining and incrementing operations is repeatedly executed at a particular temporal rate to allow for accurate network data to be relayed from the second device based on a prior incremental increase of the current bit rate.

17. The method of claim 14 further comprising:
even after the iterative performance of the receiving, determining and incrementing operations has terminated, receiving networking data regarding the communication channel from the second device, and from the received network data, identifying a change in the available resources of the communication channel;
incrementing or decrementing the current bit rate based on the identified change in the available resources of the communication channel.
18. The method of claim 14, wherein determining comprises determining that the current bit rate has reached the first ceiling bit rate.
19. The method of claim 14,
wherein the received data comprises an available bandwidth measurement data that is provided by the second device,
wherein determining comprises determining that the current bit rate has reached a second ceiling bit rate that is lower than the first ceiling bit rate and that is derived

- from the available bandwidth measurement data provided by the second device.
20. The method of claim 14, wherein the received data comprises a one way signal delay that is computed by the second device.
21. The method of claim 20, wherein the one way signal delay is computed by the second device by examining delay of audio packets sent from the first device to the second device.
22. The method of claim 14, wherein the received data comprises an available bandwidth measurement data that is computed by the second device by examining the size of video packets transmitted to the second device with the difference in time that the second device receives the packets.
23. The method of claim 14, wherein the set of instructions for receiving network data comprises a set of instructions for receiving network data in packets that are sent along a real-time transport protocol (RTP) stream and that are identified in the RTP stream by RTP Header Extensions.

* * * * *