



(51) International Patent Classification:

Not classified

(21) International Application Number:

PCT/US2021/047326

(22) International Filing Date:

24 August 2021 (24.08.2021)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

63/224,457 22 July 2021 (22.07.2021) US

(71) Applicant: **GOOGLE LLC** [US/US]; 1600 Amphitheatre Parkway, Mountain View, CA 94043 (US).

(72) Inventors: **MARCHANT, Robert**; 1600 Amphitheatre Parkway, Mountain View, CA 94043 (US). **BUTLER,**

Triona Éidín; 1600 Amphitheatre Parkway, Mountain View, CA 94043 (US). **HOLLAND, Henry John**; 1600 Amphitheatre Parkway, Mountain View, CA 94043 (US). **JONES, David Matthew**; 1600 Amphitheatre Parkway, Mountain View, CA 94043 (US). **PAWLE, Benjamin Guy Alexander**; 1600 Amphitheatre Parkway, Mountain View, CA 94043 (US). **COLVILLE, Michael**; 1600 Amphitheatre Parkway, Mountain View, CA 94043 (US). **RICKER-BY, George Joseph**; 1600 Amphitheatre Parkway, Mountain View, CA 94043 (US).


(74) Agent: **ZIDEL, Andrew T.** et al.; Botos Churchill IP Law LLP, 430 Mountain Ave., Suite 401, New Providence, NJ 07974 (US).

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

(54) Title: USING SIMPLE MASKS FOR ONLINE EXPRESSION

Fig. 3

300



(57) Abstract: The technology provides enhanced co-presence of interactive media participants without high quality video or other photo-realistic representations of the participants. A low-resolution graphical representation (318) of a participant provides real-time dynamic co-presence. Face detection captures a maximum amount of facial expression with minimum detail in order to construct the low-resolution graphical representation. A set of facial mesh data (304) is generated by the face detection to include a minimal amount of information about the participant's face per frame. The mesh data, such as facial key points, is provided to one or more user devices so that the graphical representation of the participant can be rendered in a shared app at the other device(s) (308, 804). The rendering can include generating a hull (314, 806) that delineates a perimeter of the user mask and generating a set of facial features (316, 808), in which the graphical representation is assembled by combining the hull and the set of facial features (318, 810).

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

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

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

- *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*
-

Using Simple Masks for Online Expression

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims priority to and the benefit of the filing date of Provisional Application No. 63/224,457, filed July 22, 2021, the entire disclosure of which is incorporated by reference herein.

BACKGROUND

[0002] Real-time interactive communication can take on many forms. Videoconference or other applications that employ user imagery can help provide contextual information about the participants, which can promote robust and informative communication. However, video imagery of participants can suffer from bandwidth constraints, which may be due to a poor local connection (e.g., WiFi) or issues with the communication backhaul. Regardless of the cause, bandwidth-related issues can adversely impact the users' interaction. In addition, providing photo-realistic representations of participants or the location they are communicating from can introduce other concerns.

[0003] While certain concerns relating to user-specific information may be addressed by using a static avatar, blurring the background or otherwise utilizing a selected scene as a background view, this does not alleviate problems with video of the users themselves. Approaches other than real-time video of a user may include memoji that can mimic certain facial expressions, or face filters that utilize augmented reality (AR) virtual objects to alter a person's visual appearance. However, these approaches can fail to provide useful real-time contextual information regarding a participant, including how they are interacting with other users or the application that brings the users together. They may also not be suitable for many types of interaction, especially for professional or business settings such as board meetings, client pitches, icebreakers, tele-health visits or the like.

BRIEF SUMMARY

[0004] The technology relates to methods and systems for enhanced co-presence of interactive media participants (e.g., for web apps and other applications or programs) without relying on high quality video or other photo-realistic representations of the participants. A low-resolution puppet, simulacrum or other graphical representation of a participant (a "user mask") provides real-time dynamic co-presence, which can be employed in a wide variety of video-focused, textual-focused or other types of applications. This can include videoconferencing, shared experience and gaming scenarios, document sharing and other collaborative tools, texting or chat applications, etc.

[0005] According to one aspect of the technology, a method comprises accessing, via a first computing device associated with a first user, a collaborative program configured to support multiple participants; obtaining, by the first computing device, a set of mesh data corresponding to a virtual representation of a face of a participant in the collaborative program; generating, by one or more

processors of the first computing device, a hull of a user mask, the hull delineating a perimeter of the user mask in accordance with the set of mesh data; generating by the one or more processors, a set of facial features in accordance with the set of mesh data; assembling, by the one or more processors, the user mask by combining the hull and the set of facial features; and incorporating the user mask into a graphical interface of the collaborative program.

[0006] In one example, the participant is a second user associated with a second computing device, and obtaining the set of mesh data comprises receiving the set of mesh data corresponding to the virtual representation of the face of the second user from the second computing device. In another example, generating the hull and generating the set of facial features are performed in parallel. In a further example, the method further comprises updating the user mask based on a newer set of mesh data.

[0007] In yet another example, the method further comprises performing pre-processing on the set of obtained mesh data based on a user interaction with the collaborative program. The pre-processing may include rotating point coordinates of the set of mesh data. In this case, rotating the point coordinates may cause the user mask to change orientation to indicate where the participant's focus is.

[0008] In another example, the participant is the first user, and the method further comprises generating the set of mesh data from a frame captured by a camera associated with the first computing device. In a further example, the method also comprising changing at least one of a resolution or a detail of the user mask based on a detected motion of the participant. In yet another example, the method further includes changing at least one of a resolution or a detail of the user mask based on available bandwidth for a communication connection associated with the collaborative program. And in another example, the method further comprises changing at least one of a resolution or a detail of the user mask based on computer processing usage associated with the one or more processors of the first computing device.

[0009] In addition to or complementary with the above examples and scenarios, the user mask may be configured to illustrate at least one of a facial expression or positioning of the participant's head. Upon determining that there is connectivity issue associated with the collaborative program, the method may further include locally updating the user mask without using a newer set of mesh data. The connectivity issue may indicate a loss of connection that exceeds a threshold amount of time. Generating the hull may include performing a hull concavity operation to delineate the perimeter of the user mask.

[0010] According to another aspect of the technology, a computing device is provided. The computing devices comprises memory configured to store data associated with a collaborative program, and one or more processors operatively coupled to the memory. The one or more processors are configured to: access the collaborative program, the collaborative program being configured to support multiple participants; obtain a set of mesh data corresponding to a virtual representation of a face of a

participant in the collaborative program; generate a hull of a user mask, the hull delineating a perimeter of the user mask in accordance with the set of mesh data; generate a set of facial features in accordance with the set of mesh data; assemble the user mask by combining the hull and the set of facial features; and incorporate the user mask into a graphical interface of the collaborative program.

[0011] In one example, the one or more processors are further configured to update the user mask based on a newer set of mesh data. In another example, the one or more processors are further configured to pre-process the set of obtained mesh data based on a user interaction with the collaborative program. In a further example, the user mask illustrates at least one of a facial expression or positioning of the participant's head.

[0012] In one scenario, the computing device further includes at least one camera and the participant is a user of the computing device. There, the one or more processors are further configured to generate the set of mesh data from a frame captured by the one or more cameras associated.

[0013] And in another scenario, the one or more processors are further configured to: change at least one of a resolution or a detail of the user mask based on a detected motion of the participant; change at least one of the resolution or the detail of the user mask based on available bandwidth for a communication connection associated with the collaborative program; or change at least one of the resolution or the detail of the user mask based on computer processing usage associated with the one or more processors.

BRIEF DESCRIPTION OF THE DRAWINGS

[0014] Fig. 1 illustrates an example simple mask scenario in accordance with aspects of the technology.

[0015] Figs. 2A-B illustrate a user mask approach in accordance with aspects of the technology.

[0016] Fig. 3 illustrates a user mask creation process in accordance with aspects of the technology.

[0017] Figs. 4A-B illustrate a system for use with aspects of the technology.

[0018] Figs. 5A-C illustrate different scenarios with user masks in accordance with aspects of the technology.

[0019] Figs. 6A-C illustrates a user mask presentation scenario in accordance with aspects of the technology.

[0020] Figs. 7A-B illustrate additional examples in accordance with aspects of the technology.

[0021]

[0022] Fig. 8 illustrates a method in accordance with aspects of the technology.

DETAILED DESCRIPTION

Overview

[0023] According to the technology, a face detection process is able to capture a maximum amount of facial expression with minimum detail in order to construct a “mask” of the user. Here, a facial mesh is generated at a first user device in which the facial mesh includes a minimal amount of information per frame. The facial mesh information, such as key points of the face, are provided to one or more other user devices, so that the graphical representation of the participant at the first device can be rendered in a shared app at the other device(s).

[0024] By way of example, the facial mesh information may be updated for each frame of video captured at the first user device. The information from each frame may be, e.g., on the order of 50-150 2D points at one byte per dimension, so between 100-200 bytes per frame (uncompressed). The rendered graphical representation thus illustrates real-time interaction by the participant at the first device, conveying facial expressions, overall appearance and pose with a minimal amount of transmitted information. Should quality issues degrade a video feed, the graphical representation (including the facial expressions) of the participant can remain unaffected due to its very low bandwidth requirements.

[0025] Fig. 1 illustrates an example scenario 100 employing simple masks in accordance with aspects of the technology. As shown, there is a pair of client computing devices 102a and 102b, each associated with a different user. Each client device has at least one user-facing camera 104a or 104b. In this example, the first client computing device 102a may be a laptop, while the second client computing device may be a home-related device such as a smart home assistant. However, the client computing devices may be other types of computing devices such as a desktop computer, a tablet PC, a netbook, a mobile phone, a personal digital assistant, a smartwatch or other wearable computing device, etc.

[0026] In other configurations, the user-facing camera may be separate from the computing device (e.g., a portable webcam or mobile phone on the user’s desk). In this case, the device with the camera could be paired to the computing device or configured to run a co-presence application. By way of example, a user may run the co-presence application on a desktop computing device with no camera/microphone capability, but shares their screen to show a presentation. On a secondary device, the same user would also run a version of the application, where a camera is used to capture the key facial points and audio information via a microphone.

[0027] While only two client computing devices 102a,b are shown, the technology can support three or more users and their respective devices. In this example, each client device has a common app or other program 106, such as collaborative spreadsheet. The app may be a program that is executed locally by a client device, or it may be managed remotely such as with a cloud-based app. In this example, a first graphical representation 108a (e.g., a puppet or other mask representation) is associated

with a user of the first client computing device 102a, and a second graphical representation 108b is associated with a user of the second client computing device 102b.

[0028] In one scenario, the graphical representations 108a,b are rendered locally at the respective computing devices based on facial mesh information derived from imagery captured by the cameras 104a,b. Thus, instead of transmitting a high-quality video stream or other photo-realistic representation of the user of the first client computing device that may require kilobytes of data per frame, the facial mesh information is provided for rendering in the app 106. As discussed further below, the rendered graphical representation(s) enable real-time dynamic co-presence of the app's collaborators. This can include changes to facial expressions, syncing with audio, reorienting the position of the user representation to indicate the user's focus in the app, moving the location of the user representation within the app, e.g., to show where the person is working within the app or to delineate a presenter from other participants, etc.

[0029] As shown in view 200 of Fig. 2A, when a user is looking towards the display of the client device, the camera obtains imagery of the person. Upon detecting the presence of the user, a face mesh is generated. Data points obtained from the face mesh are then used to construct a mask of the user, as shown in view 210 of Fig. 2B. This enables the system to capture maximum facial expression with minimum detail. In particular, the user mask provides the expressivity and animation of video, but free from the visual noise associated with video (e.g., due to compression artifacts, dropped packets, etc.). This, in turn, permits the user masks to be easily overlaid or otherwise incorporated into a user interface (UI) to create new collaborative experiences.

User Mask Creation

[0030] Fig. 3 illustrates an example process 300 for creating a user mask. At block 302, an image is captured by the camera of the client computing device. From this, at block 304 a face mesh is obtained. The face mesh may be generated in the following manner. In one scenario, a machine learning model is used to infer three-dimensional surface geometry of the person's face. Here, a face detection model may be employed in conjunction with a face landmark model, using only a single camera for input without a depth sensor. The face detection model is able to operate on the full image received from the camera in order to determine face location information. For instance, the face detection model may generate face boundary information (e.g., rectangles) and selected landmarks such as the centers of the eyes, ear features such as tragiions, the tip of the nose, etc.

[0031] The face landmark model relies on these landmarks to predict an approximate surface geometry (e.g., via regression). In particular, the landmarks are references for aligning the face to a selected position, for example with the eye centers along a horizontal axis of a facial rectangle. This allows for cropping of the image to omit non-facial elements. Cropping information can be generated based on face landmarks identified in a preceding frame. Once cropped (and optionally resized), the

data is input to a mesh prediction neural network, which generates a vector of landmark coordinates in three dimensions. This vector can then be mapped back to the coordinate system of the captured image from the camera. The result is a mesh that acts as a virtual representation of the user's face. The mesh data may include the vertices with selected facial regions and the center of the user's head optionally identified.

[0032] In one scenario, both image capture and face mesh generation occur in real time at the same client device (as indicated by dashed line 306). As indicated by arrow 308, the mesh data can be transmitted to the client devices of the other participants and/or used by the client device where the imagery was captured, to render the user mask. Alternatively, for cloud-based or centrally managed apps, the user mask can be rendered remotely for presentation in an app UI.

[0033] According to one aspect, the face mesh model generates landmark information, which include labelled points that constitute a specific feature of a face (e.g., the left eyebrow). This data is sent as part of the mesh of points. In one scenario, only the coordinates of key points of the face mesh are sent so that the graphical mask can be rendered at the other end. For instance, each mesh may use approximately 50-100 two-dimensional points at 1 byte per dimension, or on the order of 100-200 bytes per frame, uncompressed. This would be several orders of magnitude smaller than for a typical video call, which may operate at approximately 300 Kb/s. At 30 frames per second, that is on the order of 10,000 bytes per frame. According to one example, the model may generate a maximum of 400-600 points; however, a subset of fewer than the maximum number of points can be transmitted to save bandwidth, with a corresponding reduction in fidelity of the user mask generated at the receiving end. By way of example, the subset may be 90%, 80%, 70%, 60%, 50% or fewer of the maximum number of points.

[0034] In addition to sending the set (or subset) of the mesh of points, inference data from other models could also be sent. For instance, a model that detected sentiment based on facial expressions could be used to provide contextual information. In one implementation, a multimodal model could be used that detects all of the information that will be associated with the mask to be generated.

[0035] As shown by dashed block 310, mask generation may be performed in the following manner. Upon receipt of the key point coordinates of the face mesh, this information may be pre-processed as shown at block 312. By way of example, the pre-processing may include rotation of the key point coordinates, such as to cause the generated mask to change orientation to indicate where the user's focus is. For instance, if the user turns their head or looks towards a particular place on the display screen, this can be reflected by rotation. In conjunction with head rotation, a mouse, cursor, touchscreen and/or pointer could be used as an additional signal for focus. Pre-processing can alternatively or additionally include scaling. Other pre- or post-processing can include the application of data from other models. By way of example, if another model detected that the person had a

particular emotion e.g., was angry (or happy, sad, etc.), the face mesh point cloud could be adjusted to make the mask appear ‘angrier’ (or happier or sadder, etc.). In addition, the color and/or texture of the how the hull and features are rendered could also be altered according to the information from the other model(s).

[0036] After any pre-processing, the system generates a face “hull” as shown at block 314 and facial features as shown at block 316. In particular, at block 314 the hull or outer perimeter of the mask is drawn around the mesh points. By way of example, the hull may be generated by taking all the points of the facial features, creating a line that encircles all of the points, and performing a hull concavity operation to draw the encircling line inward, for instance until the line circumscribes the outermost points with a minimum of empty space between each point and the line. In one scenario, one or more of the mesh points may be discarded or ignored in order to create a “smooth” hull that has a generally oval or rounded appearance. At block 316, the facial features are drawn using lines and polygons between selected mesh points. These operations may be done in parallel or sequentially. Then, at block 318, the hull and facial features are assembled into the user mask. The overall process can be repeated for subsequent image frames captured by the camera.

[0037] Rather than computing and drawing the hull, the system could also triangulate the mesh and render those triangles using a 3D renderer. The hull could also be drawn more approximately, by drawing a predefined shape e.g., a circle, oval or other shape behind the mask. Another approach altogether would be to use the facial features to drive a 3D model puppet. In this case, some 3D points on a puppet model would be connected to the face mesh features, and the motion of the model would be driven by the motion of the mesh. By way of example, there could be a system linking points on the face mesh to parts of a 3D face model, so that motion in the face mesh is reflected in the 3D model. Thus, one could draw a high-resolution 3D face, with the location, rotation, and positioning of features like eyes, eyebrows and mouth driven by the face mesh data.

Example System

[0038] The user masks that are generated can be used in a wide variety of applications and scenarios, as discussed in detail below. How the user masks are generated and shared can depend on how the participants communicate with one another. One example computing architecture is shown in Figs. 4A and 4B. In particular, Figs. 4A and 4B are pictorial and functional diagrams, respectively, of an example system 400 that includes a plurality of computing devices and databases connected via a network. For instance, computing device(s) 402 may be a cloud-based server system that provides or otherwise supports one or more apps, games or other programs. Database 404 may store app/game data, user profile information, or other information with which the user masks may be employed. The server system may access the databases via network 406. Client devices may include one or more of a desktop computer 408, a laptop or tablet PC 410 and in-home devices such as smart display 412. Other client

devices may include a personal communication device such as a mobile phone or PDA 414 or a wearable device 416 such as a smart watch, etc. Another example client device is a large screen display or interactive whiteboard 418, such as might be used in a classroom, conference room, auditorium or other collaborative gathering space.

[0039] In one example, computing device 402 may include one or more server computing devices having a plurality of computing devices, e.g., a load balanced server farm or cloud computing system, that exchange information with different nodes of a network for the purpose of receiving, processing and transmitting the data to and from other computing devices. For instance, computing device 402 may include one or more server computing devices that are capable of communicating with any of the computing devices 408-418 via the network 406. This may be done as part of hosting one or more collaborative apps (e.g., a videoconferencing program, an interactive spreadsheet app or a multiplayer game) or services (e.g., a movie streaming service or interactive game show where viewers can provide comments or other feedback).

[0040] As shown in Fig. 4B, each of the computing devices 402 and 408-418 may include one or more processors, memory, data and instructions. The memory stores information accessible by the one or more processors, including instructions and data that may be executed or otherwise used by the processor(s). The memory may be of any type capable of storing information accessible by the processor(s), including a computing device-readable medium. The memory is a non-transitory medium such as a hard-drive, memory card, optical disk, solid-state, etc. Systems may include different combinations of the foregoing, whereby different portions of the instructions and data are stored on different types of media. The instructions may be any set of instructions to be executed directly (such as machine code) or indirectly (such as scripts) by the processor(s). For example, the instructions may be stored as computing device code on the computing device-readable medium. In that regard, the terms “instructions”, “modules” and “programs” may be used interchangeably herein. The instructions may be stored in object code format for direct processing by the processor, or in any other computing device language including scripts or collections of independent source code modules that are interpreted on demand or compiled in advance.

[0041] The processors may be any conventional processors, such as commercially available CPUs. Alternatively, each processor may be a dedicated device such as an ASIC, graphics processing unit (GPU), tensor processing unit (TPU) or other hardware-based processor. Although Fig. 4B functionally illustrates the processors, memory, and other elements of a given computing device as being within the same block, such devices may actually include multiple processors, computing devices, or memories that may or may not be stored within the same physical housing. Similarly, the memory may be a hard drive or other storage media located in a housing different from that of the processor(s), for instance in a cloud computing system of server 402. Accordingly, references to a processor or computing device

will be understood to include references to a collection of processors or computing devices or memories that may or may not operate in parallel.

[0042] The computing devices may include all of the components normally used in connection with a computing device such as the processor and memory described above as well as a user interface subsystem for receiving input from a user and presenting information to the user (e.g., text, imagery and/or other graphical elements). The user interface subsystem may include one or more user inputs (e.g., at least one front (user) facing camera, a mouse, keyboard, touch screen and/or microphone) and one or more display devices that is operable to display information (e.g., text, imagery and/or other graphical elements). Other output devices, such as speaker(s) may also provide information to users.

[0043] The user-related computing devices (e.g., 408-418) may communicate with a back-end computing system (e.g., server 402) via one or more networks, such as network 406. The user-related computing devices may also communicate with one another without also communicating with a back-end computing system. The network 406, and intervening nodes, may include various configurations and protocols including short range communication protocols such as Bluetooth™, Bluetooth LE™, the Internet, World Wide Web, intranets, virtual private networks, wide area networks, local networks, private networks using communication protocols proprietary to one or more companies, Ethernet, WiFi and HTTP, and various combinations of the foregoing. Such communication may be facilitated by any device capable of transmitting data to and from other computing devices, such as modems and wireless interfaces.

Exemplary Applications and Scenarios

[0044] How the user masks of participants are displayed may depend on the type of app, game or other program, what a given participant is doing at a particular point in time, the number of participants, the size of the display screen and/or other factors. This is explored in the following example scenarios.

[0045] Fig. 5A illustrates a videoconference-type scenario 500, in which different participants either have their own client computing devices 502 and 504, or are sharing a device such as an interactive whiteboard 506. While users may conventionally interact with full video imagery from their devices' video cameras, in this scenario one or more of the participants may have bandwidth constraints (e.g., poor WiFi connectivity or spotty cellular coverage). Or a participant may not want to have themselves or their personal workspace appear on camera. As shown on the displays of devices 502 and 504, each user mask may be displayed in a specific panel or section of the UI. Alternatively, as shown on the display of whiteboard 506, the participants that are in the same room as the whiteboard may be grouped together in the same panel or section 508 of the UI. By way of example, the UI could vary the thickness or color of the mask hull to indicate that a particular person is speaking (or taking some action). Alternatively or additionally, more mesh points could be used for a person that is speaking

than a person in the audience. This could be done by reducing the mesh points transmitted for passive participants, e.g., to reduce bandwidth, and/or to transmit more mesh points for an active contributor.

[0046] In another scenario, participants who are not “active” (e.g., for more than a threshold amount of time such as 2-5 minutes, or more or less) could have their masks generated locally by the receiving computer (for example without transmitting any face mesh data), to stop them from “going static”. Here, the masks would effectively be switched to a “self-driving” mode when they don't have a lot of input from their associated user. In a further scenario, the “self-driving” mode may be activated when a participant's connection is poor, and little to no face data has been received for some threshold period of time (e.g., on the order of 1-5 seconds or more). Here, the participant can continue to have their attention communicated even when the connection is unreliable.

[0047] Furthermore, the user masks could be augmented with symbols, emoji, etc. to add contextual information such as status, activity or the like. Here, the mesh could be distorted or have transformations applied to it to convey information, such as that the user has looked away from the camera, so there is not a clear mesh to transmit. In this case, the recipient computer(s) could shrink the displayed user mask of the other participant to indicate that state.

[0048] Fig. 5B illustrates a texting-type scenario 520, in which the participants are able to chat via textual input. Here, user interface 522 includes a first pane 524 that indicates different users that are in a group or are otherwise able to chat with one another. In this example, a second pane 526 is shown along the top of the UI, indicating the user's name and how many other people are actively in the chat. And a third pane 528 presents the textual communication between the participants. As shown, the participants are able to easily jump into a more robust interaction with user masks providing contextual cues that are not available when just using text. Here, any participant can select the Chat Call icon 530 to initiate visual communication via user masks. In this scenario, audio may be optionally provided to complement the user masks. This approach can enhance the communication without the formality and direct focus of a video call. It enables participants to chat as normal while reading social cues that would otherwise be missing from the text communication.

[0049] Fig. 5C illustrates a collaborative document editing scenario 540, in which a document 542 in UI 544 is being shared among different users (here, Jo, Sunil and Carli). While the document 542 is shown as a report, it may be any type of document (e.g., a spreadsheet, presentation, etc.). In this type of scenario, the masks can be arranged in the UI to create awareness of a participant's state of mind, mood and/or area of concentration. This can promote better riffing and enable polite interruptions in a bandwidth-efficient manner. Collaborators are able to focus on the content of the document, and not on the other people per se.

[0050] By way of example, the mask locations show where different people are looking in the document, and a tilt of the mask can show where they are typing or selecting content. For instance, as

shown by dotted line 546, Jo's user mask is tilted to face towards the word "istance", which is a typographical error. And Sunil's mask is facing towards the bolded text "exponential" as they type that word into the document as shown by dashed line 548. Here the cursor may be augmented by the user mask, such as to provide more context to the other collaborators as to why the document is being modified. Or the cursor can act as an anchoring signal for the mask, so that rotation or other movement of the mask is made in relation to the position of the cursor. Also shown in this example is Carli's mask. Here, when Carli begins to speak with the other participants, mask 550 turns face-on (not angled toward the textual content) so that Carli's facial expression can be more clearly illustrated along with any accompanying audio.

[0051] Figs. 6A-C illustrate another scenario 600, in which audience reaction can help shape a remote presentation. In this example, Sunil is giving a presentation about sales projections. Here, the user masks of other participants can provide visual feedback to the presenter, allowing Sunil to tailor the presentation to the audience. This is done without distracting video feeds from the audience (or Sunil) competing with the content that is the focus of the presentation. As seen in Fig. 6A, the positions of the user masks around the UI imply social cues. For instance, Sunil's mask is adjacent to the slide of the presentation that is currently being displayed (e.g., along a first pane 602 of the UI), which indicates that Sunil is the presenter, while the other masks are positioned separately in a different pane 604 of the UI (e.g., an "auditorium" section) to indicate that the other participants are the audience.

[0052] As seen in Fig. 6B, when user attention from audience members focuses on a particular part of the slide, their user mask moves toward that part of the slide. In view 610 of Fig. 6B, assume that Jo and Shannon are focused on the December monthly forecast portion of the slide. Here, not only can Sunil see that Jo and Shannon are looking at the December portion, but can also see their real-time facial expressions. Based on this, the presenter can quickly jump to the relevant slide that discusses that area of interest in more detail. For instance, as shown in view 620 of Fig. 6C, a pie chart slide shows the relative sales projections for widgets of various sizes. As seen here, while Sunil is still presenting, it can be easily discerned that the other participants are focused on specific portions of the chart, including their expressions. Thus, the presenter is able to quickly and effectively tailor the presentation based on visual cues from other participants. It will be appreciated that, particularly in situations where the audience has many members, the bandwidth savings associated with presenting an array of user masks in comparison with an array of video feeds for the members, are significant. There is also an efficiency in terms of the user of the screen real-estate; useful information for each audience member, via a spatially efficient user mask, can be readily conveyed to the presenter to guide their interaction with the presentation application.

[0053] View 700 of Fig. 7A shows another example application for a remote movie night with friends. Here, the viewers masks are overlaid in the display screen to provide a shared experience. For

instance, the friends would each stream the same movie (or tv show, sporting event, etc.) and would be able to see everyone else's reactions via their masks' expressions. Similarly, view 710 of Fig. 7B shows an example e-sports scenario where friends or other gamers can be involved in a player's game. In this case, the player may have a heads-up display with their friends' masks shown along a perimeter of the screen. This allows the player to see their in-game reactions to how he/she is performing.

[0054] These are just a few examples of how user masks may be employed to enrich the interaction of users in a shared space, without the distraction or other downsides to employing full motion video of each person. Of course, this would not prevent some users from having a video of them presented to other participants. However, in one scenario should there be a disruption in a WiFi signal or bandwidth issue with the connection, then the system could gracefully fall back to presenting the user mask instead of full motion video. For instance, should a problem with the video be detected such as falling below a threshold quality level, then the face mesh for a user can be automatically generated by the client computing device. Alternatively, the participant may elect to switch from full motion video to the user mask, for instance by selecting a "mask" button in the app, game or other program being used. And as noted above with regard to Fig. 4B, a participant may enhance a text-based app by introducing user masks (with or without accompanying audio). In one scenario, the switch can be instantaneous (not perceptible to the participant). This switch could be accompanied by a visual transition, such as a fade from the full motion video of the person's face to the user mask. Here, the tracking model could continuously run so that there is always a face mesh ready to use.

[0055] Depending on the number of participants or other factors, it may be desirable to personalize or otherwise choose how a person's mask will be displayed. For instance, a person may be able to customize the appearance of their mask. By way of example, tunable features may include adjusting the refresh (or frame) rate, the mask color(s), resolution, size, line thickness or the like. Such modifications may be associated with a particular app (e.g., one person may have different masks for different apps), with a particular device or physical location (e.g., different appearance depending on whether the user is working from home on their laptop or mobile phone, is at work at their desk or in a conference room, etc.)

[0056] Fig. 8 illustrates a method 800 in accordance with aspects of the technology. At block 802, the method includes accessing, via a first computing device associated with a first user, a collaborative program configured to support multiple participants. At block 804, the method includes obtaining, by the first computing device, a set of mesh data corresponding to a virtual representation of a face of a participant in the collaborative program. At block 806, the method includes generating, by one or more processors of the first computing device, a hull of a user mask, the hull delineating a perimeter of the user mask in accordance with the set of mesh data. At block 808, the method includes generating by the one or more processors, a set of facial features in accordance with the set of mesh data. At block

810, the method includes assembling, by the one or more processors, the user mask by combining the hull and the set of facial features. At block 812, the method includes incorporating the user mask into a graphical interface of the collaborative program.

[0057] Although the technology herein has been described with reference to particular embodiments, it is to be understood that these embodiments are merely illustrative of the principles and applications of the present technology. It is therefore to be understood that numerous modifications may be made to the illustrative embodiments and that other arrangements may be devised without departing from the spirit and scope of the present technology as defined by the appended claims.

[0058] The user mask technology can be easily employed in all manner of apps, games and other programs, providing rich contextual information to other participants in real-time using mesh data derived from a user's face. This is done at a small fraction of the bandwidth that would be required from conventional full video imagery. By way of example, while aspects of the technology enable collaboration for documents and other text-based applications (e.g., chatting), the technology is applicable in many other contexts. For instance, friends may share a common activity such as watching a movie or gaming. Visual cues from a virtual audience can help the presented focus their attention on particular details or rapidly change slides to more relevant content.

CLAIMS

1. A method, comprising:
 - accessing, via a first computing device associated with a first user, a collaborative program configured to support multiple participants;
 - obtaining, by the first computing device, a set of mesh data corresponding to a virtual representation of a face of a participant in the collaborative program;
 - generating, by one or more processors of the first computing device, a hull of a user mask, the hull delineating a perimeter of the user mask in accordance with the set of mesh data;
 - generating by the one or more processors, a set of facial features in accordance with the set of mesh data;
 - assembling, by the one or more processors, the user mask by combining the hull and the set of facial features; and
 - incorporating the user mask into a graphical interface of the collaborative program.
2. The method of claim 1, wherein:
 - the participant is a second user associated with a second computing device; and
 - obtaining the set of mesh data comprises receiving the set of mesh data corresponding to the virtual representation of the face of the second user from the second computing device.
3. The method of claim 1 or 2, wherein generating the hull and generating the set of facial features are performed in parallel.
4. The method of any one of claims 1, 2 or 3, wherein the method further comprises updating the user mask based on a newer set of mesh data.
5. The method of any one of claims 1-4, wherein the method further comprises performing pre-processing on the set of obtained mesh data based on a user interaction with the collaborative program.
6. The method of claim 5, wherein the pre-processing includes rotating point coordinates of the set of mesh data.
7. The method of claim 5 or 6, wherein rotating the point coordinates causes the user mask to change orientation to indicate where the participant's focus is.

8. The method of any one of claims 1-7, wherein the user mask illustrates at least one of a facial expression or positioning of the participant's head.

9. The method of any one of claims 1-8, wherein the participant is the first user, and the method further comprises generating the set of mesh data from a frame captured by a camera associated with the first computing device.

10. The method of any one of claims 1-9, further comprising changing at least one of a resolution or a detail of the user mask based on a detected motion of the participant.

11. The method of any one of claims 1-10, further comprising changing at least one of a resolution or a detail of the user mask based on available bandwidth for a communication connection associated with the collaborative program.

12. The method of any one of claims 1-11, further comprising changing at least one of a resolution or a detail of the user mask based on computer processing usage associated with the one or more processors of the first computing device.

13. The method of any one of claims 1-12, wherein upon determining that there is connectivity issue associated with the collaborative program, the method further includes locally updating the user mask without using a newer set of mesh data.

14. The method of claim 13, wherein the connectivity issue is a loss of connection that exceeds a threshold amount of time.

15. The method of any one of claims 1-14, wherein generating the hull includes performing a hull concavity operation to delineate the perimeter of the user mask.

16. A computing device, comprising:
memory configured to store data associated with a collaborative program; and
one or more processors operatively coupled to the memory, the one or more processors being configured to:
access the collaborative program, the collaborative program being configured to support multiple participants;

obtain a set of mesh data corresponding to a virtual representation of a face of a participant in the collaborative program;

generate a hull of a user mask, the hull delineating a perimeter of the user mask in accordance with the set of mesh data;

generate a set of facial features in accordance with the set of mesh data;

assemble the user mask by combining the hull and the set of facial features; and

incorporate the user mask into a graphical interface of the collaborative program.

17. The computing device of claim 16, wherein the one or more processors are further configured to update the user mask based on a newer set of mesh data.

18. The computing device of claim 16 or 17, wherein the one or more processors are further configured to pre-process the set of obtained mesh data based on a user interaction with the collaborative program.

19. The computing device of any one of claims 16-18, wherein the user mask illustrates at least one of a facial expression or positioning of the participant's head.

20. The computing device of any one of claims 16-19, wherein:

the computing device further includes at least one camera;

the participant is a user of the computing device; and

the one or more processors are further configured to generate the set of mesh data from a frame captured by the one or more cameras associated.

21. The computing device of any one of claims 16-20, wherein the one or more processors are further configured to:

change at least one of a resolution or a detail of the user mask based on a detected motion of the participant;

change at least one of the resolution or the detail of the user mask based on available bandwidth for a communication connection associated with the collaborative program; or

change at least one of the resolution or the detail of the user mask based on computer processing usage associated with the one or more processors.

Fig. 1
100




Fig. 2A
200







Fig. 2B
210











Fig. 4B






Fig. 5C
540




Fig. 6A
600




Fig. 6B
610





Fig. 6C
620



12/14

Fig. 7A
700



Fig. 7B
720

