



- (51) International Patent Classification:
H04N 7/15 (2006.01) *H04L 29/06* (2006.01)
- (21) International Application Number:
PCT/US2017/044062
- (22) International Filing Date:
27 July 2017 (27.07.2017)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
15/230,202 05 August 2016 (05.08.2016) US
- (71) Applicant: MICROSOFT TECHNOLOGY LICENSING, LLC [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).
- (72) Inventor: FAULKNER, Jason Thomas; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US).

- (74) Agent: MINHAS, Sandip S. et al.; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,

(54) Title: COMMUNICATION EVENT

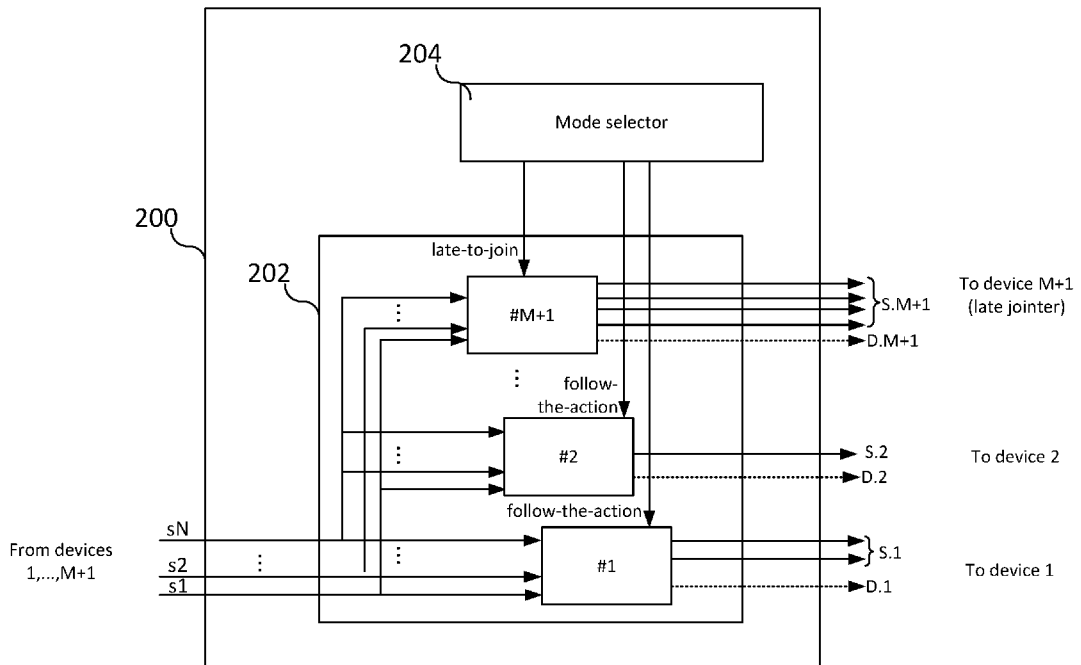


FIG. 2

(57) Abstract: An established communication event is conducted between existing participants by transmitting and receiving video streams via a communication network between their user devices. A user interface of a new participant is controlled to operate in a late-to-join mode initially if that participant satisfies a late-to-join condition. In response to expiry of an initial duration, the user interface automatically switches to an active speaker mode instead.

WO 2018/026615 A1

TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

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

Published:

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

COMMUNICATION EVENT

Technical Field

[0001] The present invention relates to a communication event conducted by
5 transmitting and receiving video streams via a communication network between user
devices operated by participants of the communication event, and displaying video images
derived from those video streams at those user devices.

Background

[0002] Voice over internet protocol (“VoIP”) communication systems allow a user
10 of a device to conduct a communication event, such as a call, via a packet-based
communication network such as the Internet. To use VoIP, the user installs and execute
client software on a user device. The client software provides the VoIP connections as well
as other functions such as registration and authentication. In a call conducted using the
clients, audio data is transmitted and received between the clients so that users participating
15 in the call can hear each other.

[0003] In addition to voice communication, the client may also provide video calling
functionality, such that video streams carrying video images of the participants can be
transmitted and received between the clients via the network. The participants are able to
view the video images of the other participants in addition hearing their voices. This
20 provides a more natural form of communication between the participants, as facial
expressions are also communicated, thereby making video calls more comparable to a face-
to-face conversation. The video images can be call video images i.e. images of the users
themselves captured at their devices, but they can also be other types of video image (that
is, other types of video content) in order to implement other communication event
25 modalities, such as screen sharing, document sharing, whiteboard sharing etc. For example,
at least one of the video streams may carry a video image of content currently being
displayed one of the participant's displays so that it is visible to the other participants.

Summary

[0004] An established communication event is conducted between existing
30 participants by transmitting and receiving video streams via a communication network
between their user devices. A user interface of a new participant is controlled to operate in
a late-to-join mode initially if that participant satisfies a late-to-join condition. In response
to expiry of an initial duration, the user interface automatically switches to an active speaker
mode instead.

[0005] A first aspect of the present invention is directed to a method of controlling a user interface of a user device operated by a new participant of an established communication event, the established communication event being conducted between a plurality of existing participants by transmitting and receiving video streams via a communication network between user devices operated by those existing participants, the method comprising steps of:

[0006] receiving a notification of the new participant joining the established communication event;

[0007] determining, based on the received notification, whether the new participant satisfies a late-to-join condition; and

[0008] if the late-to-join condition is determined to be satisfied: controlling the user interface to operate in a late-to-join mode for an initial duration, and in response to expiry of the initial duration, controlling the user interface to operate in an active speaker mode instead (i.e. switching from the late-to-join mode to the active speaker mode).

[0009] controlling the user interface to operate in the active speaker mode comprises selecting for the new participant, based on audio data generated at the user device of an actively speaking one of the participants: one of the video streams of the actively speaking participant, or an avatar image representing the actively speaking participant, thereby causing the user interface to display the selected avatar image or a video image derived from the selected stream.

[0010] The terms "active-speaker mode" and "follow-the-action mode" are used interchangeably herein to refer to a mode in which the displaying of video images or avatars is driven primarily, but not necessarily exclusively, on active speaker identification, implemented by applying speech detection to audio data of the participants.

[0011] Controlling the user interface to operate in the late-to-join mode may comprise one of the video streams of a non-speaking participant or an avatar image of the non-speaking participant, thereby causing the user interface to display the selected avatar image or a video image derived from the selected stream of the non-speaking participant.

[0012] A video image can be a "moving image", i.e. a sequence of frames (i.e. static images) to be played out in quick succession to give the impression of movement during a communication event, or it can be a image still from a moving image (that is, a static image derived from a moving image). Unless otherwise indicated, any references to "an image" below denote a video image in this sense. The phrase "displaying a video stream" mean displaying a video image derived from that video stream. A video image of the

communication event means a video image derived from one of the video streams of the communication event. Still images are potential views that can be displayed as part of a late to join experience. For example, a screen capture of a recent high activity view content sharing.

5 **[0013]** The video streams can carry video images of the users themselves ("call video image") but the present invention is not limited to this; at least one may carry a video image of another type of content, such as a shared content video image ("shared content media image"). In this respect, whilst a "call video image" can be an image of one or more user's, the invention is not limited to this – at least one of the video images could be of
10 shared content. For example, in-app document sharing/editing, or screen sharing. From the perspective of the new participant, the shared content appears to an extent as a "participant" in the communication event in its own right, referred to as "content-as-a-participant", although it is in fact video content supplied by one of the actual participants (i.e. one of the user's participating in the communication event). For the avoidance of any doubt, it is noted
15 that the term "participant" when used in the Claims means a user participating in the communication event.

[0014] For example, one or more dominant activity image stream(s) plus one or more content media streams could be displayed.

[0015] As will be apparent, "initial" in this context is from the perspective of the
20 new user, and is a time interval commencing when he has joined (it does not refer to the very start of the communication event, which is before he has joined).

[0016] In embodiments, controlling the user interface to operate in the late-to-join mode may comprise selecting for the new participant one or more of said video streams and/or one or more avatar images independently of any audio data generated at the user
25 devices, thereby causing the user interface to display the selected avatar images or one or more video images derived from the selected video streams for at least part of the initial duration.

[0017] Alternatively, or in addition, no more than a maximum number of video images derived from the video streams may be displayed by the user interface at any one
30 time when operating in the active speaker mode, and more than that maximum number of video images may be displayed simultaneously by the user interface operating in the late-to-join mode at least once during the initial duration.

[0018] Alternatively, or in addition, controlling the user interface to operate in the late-to-join mode may comprise controlling the user interface to display over the course of

the initial duration, for every one of the existing participants, a respective video image derived from one of the video streams or a respective avatar image (not necessarily simultaneously).

5 [0019] In embodiments, the late-to-join condition may be based on at least one of an elapsed time of the communication event at a join time of the new participant. Alternatively, or in addition, it may be based on a number of existing participants of the communication event at the join time.

10 [0020] In other words, a late joiner may be classes as such because they have joined at a late stage in time, when there are a large number of existing participants already, or both.

[0021] The method may further comprise a step of controlling the user interface to operate in the active speaker mode immediately if the late-to-join condition is determined not to be satisfied. "Immediately" means the active speaker mode is the first selected mode i.e. the late-join mode is not selected initially.

15 [0022] The steps may be implemented at a relay node other than the new participant's user device.

[0023] For example, the controlling step may be implemented by transmitting from the relay node via the network to the new participant's user device any of the video streams needed by the user interface operating in the active speaker mode to display the image (and only that/those video stream(s)). No video streams are transmitted for as long as only an avatar image(s) is being displayed by the user interface.

[0024] The relay node may be a relay server, or the user device operated by one of the existing participants.

[0025] Wherein the steps are implemented by the new participant's user device.

25 [0026] For example, the user device may subscribe to the video stream of the actively speaking participant when the user interface is operating in the active-speaker mode, wherein the user device receives the subscribed-to stream from a relay node via the network.

[0027] In the late-to-join mode, the user device may subscribe to all (or some) of the video streams, wherein the user device receives the subscribed-to video streams from the relay node via the network. For example, at least one video stream of a non-speaking participant.

[0028] In the late-to-join mode, the user device may subscribe to at least one of the video streams of a non-speaking participant wherein the user device receives the subscribed-to video stream from the relay node via the network.

5 [0029] The user device can subscribe to a given stream(s) by transmitting at least one active speaker subscription message to the relay node.

[0030] At least one of the displayed images may be a call video image, derived from one of the video streams, of the actively speaking participant and is selected for the subsequent duration based on the active speaker's audio data.

10 [0031] At least one video images may be displayed by the user interface operating in the late-to-join mode during the initial duration is a call video image of a participant who is not actively speaking at that time.

[0032] Different video or avatar images may be displayed at different times during the initial duration by the user interface operating in the late-to-join mode.

15 [0033] The user device operated by at least one of the plurality existing participants of the communication event may, when the new participant joins, not be transmitting a video stream, and an avatar image representing the additional participant may be displayed at the new participant's user device by the user interface operating in the late-to-join mode for at least part of the initial duration, which is not displayed by the user interface when operating in the active speaker mode.

20 [0034] At least one of the video streams may carry a modality indicator and the method comprises detecting the modality indicator in that stream; wherein a video image derived from that video stream is displayed by the user interface operating in the late-to-join mode for at least part of the initial duration; and wherein the detection of the modality indicator prevents any video image derived from that video stream from being displayed by
25 the user interface when operating in the active speaker mode.

[0035] A second aspect of the invention is directed to a network node for controlling a user interface displayed to a new participant of an established communication event, the established communication event being conducted between a plurality of existing participants by transmitting and receiving video streams via a communication network
30 between user devices operated by those existing participants, the network node comprising:

[0036] an input configured to receive a notification of a new participant joining the established communication event;

[0037] a controller configured to implement steps of:

[0038] determining, based on the received notification, whether the new participant satisfies a late-to-join condition; and

[0039] if the late-to-join condition is determined to be satisfied: controlling the user interface to operate in a late-to-join mode for an initial duration, and in response to expiry
5 of the initial duration, controlling the user interface to operate in an active speaker mode instead;

[0040] wherein controlling the user interface to operate in the active speaker mode comprises selecting for the new participant, based on audio data generated at the user device of an actively speaking one of the participants: one of the video streams of the actively
10 speaking participant, or an avatar image representing the actively speaking participant, thereby causing the user interface to display the selected avatar image or a video image derived from the selected stream.

[0041] The network node can be a relay node or a user device operated by the new participant, for example:

[0042] A fourth aspect of the present invention is directed to a method of managing
15 an established communication event being conducted between a plurality of existing participants by transmitting and receiving video streams via a communication network between user devices operated by those existing participants, the method comprising steps of:

[0043] receiving a notification of a new participant joining the established
20 communication event;

[0044] determining for a display area a minimum number of images of the communication event to be displayed to the new participant for an initial duration of the communication event ("late-to-join phase");

[0045] based on that determining step, controlling a display of a user device
25 available to the new participant to display, in a display area of the display, for the initial duration of the communication event, at least two images of the communication event, each of those images being a video image derived from one of the video streams or an avatar image representing one of the participants, such that the total number of images of the communication event displayed at that user device in said display area is no less than the
30 minimum number determined for the initial duration at all times during the initial duration;

[0046] determining for the display area a maximum number of video images to be displayed to the new participant for a subsequent duration of the communication event;

[0047] controlling the display of the new participant's user device to display, in said display area, for a subsequent duration of the communication event ("follow-the-action phase"), at least one image of the communication event, that image being a video image derived from at least one of the video streams or an avatar image representing one of the participants, such that the total number of video images of the communication event displayed at that user device in said display area does not exceed the maximum number determined for the subsequent duration at all times during the subsequent duration.

[0048] An image of the communication event means either a video image derived from one of the video streams or an avatar image representing one of the participants (e.g. an audio-only participant, who is not transmitting a video stream).

[0049] In other words, in accordance with the first aspect of the present invention:

[0050] for the initial duration, a **lower** limit is imposed on the **total** number of displayed images of the communication event (video images and/or avatar images), such that the total number of images in the display area does not fall below this at any time during the initial duration;

[0051] for the subsequent duration, an **upper** limit is imposed on the number of displayed **video** images specifically (i.e. specifically images that are derived from the video streams, as opposed to avatar images), such that the number of **video** images does not exceed this upper limit at any time (though any number of **avatar** images may be displayed in the display area in addition to this).

[0052] Preferably the maximum number determined for the subsequent duration is less than the minimum number determined for the initial duration.

[0053] For example, in the embodiments described below, for the initial duration (late-to-join phase), no less than three images of the communication event are displayed for this duration, and these can be video images, avatar images or a combination of both; for the subsequent duration (follow-the-action phase), at most two **video** images are displayed in the display area for this duration, plus any number of **avatar** on top of this.

[0054] The display of the new participant's user device has a total available display area. The display area refer to above can be all of the total display area, or just a portion. In any event, the display area is the same for the initial and subsequent duration.

[0055] In embodiments of the present invention, the steps may be implemented at a relay node other than the new participant's user device. For example, the controlling steps may be implemented by transmitting from the relay node via the network to the new participant's user device, in the first and second controlling steps, any of the video streams

needed to display the at least two and the at least one images of the communication event respectively (and only those video streams).

[0056] For example, the relay node may be a relay server, or it may be the user device operated by one of the existing participants.

5 [0057] Alternatively, the steps may be implemented by the new participant's user device.

[0058] The at least one image displayed in the subsequent duration may be a call video image of at least one of the participants derived from one of the video streams. For example, the call video image may be a video image of an actively speaking one of the participants and is selected for the subsequent duration based on associated voice data.

10 [0059] Alternatively, or in addition, the at least one image may be a video image selected for the subsequent duration based on its visual content.

[0060] Over the initial duration, video images derived from every one of the video streams may be displayed at the user device in said display area.

15 [0061] At least one of the images displayed in said display area during the initial duration may be a call video image, derived from one of the video streams, of a participant who is not actively speaking at that time.

[0062] The steps may further comprise determining a maximum number of images of the communication event to be displayed for the initial duration, wherein the new participant's user device is controlled in the first controlling step such that the number of images of the communication event displayed by that user device in said display area remains between the determined minimum and maximum numbers at all times during the initial duration.

20 [0063] Different images of the communication event may be displayed at different times during the initial duration.

[0064] The steps may comprise determining a minimum number of images of the communication event to be displayed for the subsequent duration, wherein the new participant's user device is controlled in the second controlling step such that the total number of images of the communication event displayed in said display area is no less than the minimum number determined for the subsequent duration at all times during the subsequent duration.

30 [0065] The minimum number may be determined for the initial duration by selecting for said display area one of a set of predetermined display layouts having that number of display regions.

[0066] The number of said participants throughout the subsequent duration may remain at least as great as throughout the initial duration.

[0067] The user device operated by at least one of the existing participants may, when the new participant joins, not be transmitting a video stream (audio-only participant),
5 wherein an avatar image representing the additional participant is displayed at the new participant's user device in said display area for at least part of the initial duration, which is not displayed in said display area at any time during the subsequent duration (though it may or may not be displayed outside of the available display area if said display area is less than the total display area). Alternatively, it may be displayed in the subsequent duration, such
10 that is occupied a smaller portion of the display area in the subsequent duration than the initial duration.

[0068] At least one of the video streams may carry a modality indicator and the method comprises detecting the modality indicator in that stream. A video image derived from that video stream may be displayed in said display area at the new participant's user
15 device for at least part of the initial duration. The detection of the modality indicator may prevent any video image derived from that video stream from being displayed at the new participant's user device in said display area at any time during the subsequent duration (though it may or may not be displayed outside of the available display area if said display area is less than the total display area).

[0069] At least one of the images of the communication event displayed in the initial
20 duration may be a first avatar image representing one of the participants, the at least one image of the communication event displayed in the subsequent duration may be a second avatar image representing that same participant, wherein the second avatar image occupied a smaller portion of the display area than the first avatar image. That is, the avatar
25 representation of that participant may be more prominent in the initial duration than in the subsequent duration.

[0070] A fifth aspect of the present invention is directed to a network node for managing an established communication event being conducted between a plurality of existing participants by transmitting and receiving video streams via a communication
30 network between user devices operated by those existing participants, the network node comprising: an input configured to receive a notification of a new participant joining the established communication event; a controller configured to: determine for a display area (i) a minimum number of images of the communication event to be displayed to the new participant for an initial duration of the communication event, and (ii) a maximum number

of video images to be displayed to the new participant for a subsequent duration of the communication event; and a stream selector configured to, based on the determined minimum number: (i) control a display of a user device available to the new participant to display, in a display area of the display, for the initial duration of the communication event, at least two images of the communication event, each of those images being a video image derived from one of the video streams or an avatar image representing one of the participants, such that the total number of images of the communication event displayed at that user device in said display area is no less than the minimum number determined for the initial duration at all times during the initial duration, and (ii) based on the determined maximum number, control the display of the new participant's user device to display, in said display area of the display, for a subsequent duration of the communication event, at least one image of the communication event, that image being a video image derived from one of the video streams or an avatar image representing one of the participants, such that the total number of video images of the communication event displayed at that user device in said display area does not exceed the maximum number determined for the subsequent duration and the display area at all times during the subsequent duration.

[0071] The network node may be a relay node or a user device, for example.

[0072] In embodiments, the network node can be configured to implement any of the method steps disclosed herein.

[0073] For the initial duration, at least one selectable user interface element may be displayed outside of said display area, which is displayed within said display area for the subsequent duration (e.g. in a portion of the available displayed area rendered available by the reduction in the number of displayed video images).

[0074] Said display area can be a total available display area of the display. Alternatively, it may be only part of the total available display area.

[0075] To provide what is referred to herein as a "staged-content" view, said display area in which the selected video stream(s) may be only a part of the total available display area. A shared content video image of the communication event (generated locally at the participant's own user device, or derived from another of the video streams of another of the participants) so as to occupy another, larger part of the total display area (i.e. having a greater area).

[0076] A self-video image of the new participant, generated locally at his device, may also be displayed on the display outside of said display area in which the selected video stream(s) are displayed.

[0077] Note "notification" in this context refers to a signal informing a controlling entity that is implementing the method of the new participant, such as a system notification, as opposed to a visual notification that is displayed to one of the participants. However, in some embodiments, a visual notification of the new participant may also be displayed at the user devices of the existing participants.

[0078] Another aspect of the present invention is directed to a computer program product comprising code stored on a computer readable storage medium and configured when executed to implement any of the functionality disclosed herein. The computer readable storage medium can for example be formed of one or more electronic storage devices (such as electronic, magnetic and/or optical storage), as distinct from a signal bearing medium.

[0079] User or system tagged activity can also be referenced for late to join priority. High activity content sharing activity can trigger priority for late to join grid view population for follow the action and attendee awareness on join.

[0080] A sixth aspect of the present invention is directed to a method of managing an established communication event being conducted between a plurality of existing participants by transmitting and receiving video streams via a communication network between user devices operated by those existing participants, the method comprising steps of: receiving a notification of a new participant joining the established communication event; determining a minimum number of video images of the communication event to be displayed to the new participant for an initial duration of the communication event; and based on the determining step:

[0081] controlling a display of a user device available to the new participant to display, in a display area of the display, for the initial duration of the communication event, at least two video images derived from at least two of the video streams, such that the total number of video images of the communication event displayed at that user device in said display area is no less than the minimum number determined for the initial duration at all times during the initial duration, and

[0082] controlling the display of the new participant's user device to display, in said display area of the display, for a subsequent duration of the communication event, at least one video image derived from at least one of the video streams, such that the total number of video images of the communication event displayed at that user device in said display area is less than the minimum number determined for the initial duration at all times during the subsequent duration.

Brief Description of Figures

[0083] For a better understanding of the present invention, and to show how embodiments of the same can be carried into effect, reference is made to the following figures in which:

- 5 • Figure 1 shows a schematic block diagram of a communication system;
- Figure 2 shows a functional block diagram of a relay node;
- Figure 3 shows a set of selectable predetermined call video layouts;
- Figures 4A to 4D show various illustrations of a graphical user interface (GUI) of a client operating in different call video display modes; and
- 10 • Figure 5 shows a client GUI which includes a visual notification of a late joiner;
- Figure 6A shows a first example of a client GUI;
- Figure 6B shows a second example of a client GUI;
- Figures 7A and 7B show a staged-content client GUI, operating in late-to-join and follow-the-action modes respectively.
- 15

Details Description of Preferred Embodiments

[0084] In a communication event conducted between users (participants) via a communication network, such as a video call, it may not be desirable to display, to a given participant, all of the video images of the communication event of the other participants all of the time.

20

[0085] In the examples described below, the video images are call video images i.e. video images of the participants themselves. However, as noted, the invention is not limited to this and all relevant description below pertaining to any call video image applied equally to other types of video images, such as shared content.

[0086] For example, where only a few of the users are active, i.e. doing something that may be of interest to other users in the videoconference, and/or if there is a very large number of users on the call. For example, only one or two of the users may be actively speaking, reacting, or doing users at a given time during the call. For example, a single speaking participant may be leading the call at a particular time, or two of the participants may be engaged in conversation at a particular time (whilst the others just listen).

25

30

[0087] In order to provide a more engaging user experience in such circumstances, it can be beneficial to intelligently limit the number of video images displayed to a given participant or limited number of participants, so as to focus the participant's attention on the

"current action" (i.e. those participant(s) who are currently contributing the most to the communication event) based on activity exhibited by the participants during the call.

[0088] This intelligent selection and displaying of a limited number of video streams based on participant activity is referred to herein as "follow-the-action storytelling", and
5 guides the consuming participants with the group activity and group response as the communication event proceeds. In the examples described below, when a follow-the-action mode is selected for a given user, the number of video images displayed to that user is limited to one or more, which are intelligently selected in the manner described below. This in turn provides what is referred to herein as a "storied experience view" for each participant,
10 which harnesses the power of video and storytelling to transform a meeting experience (i.e. video call) beyond active speaking via a more engaging and life-like meet up experience.

[0089] A set of one of more of the video streams can be selected for each of the participants independently, so as to provide that participant with a unique storied experience tailored to him. For example, the video stream(s) for a participant who is currently speaking
15 may be different from that or those selected for a participant who is not speaking.

[0090] The video streams to-be displayed in follow-the action storytelling can for example be selected based on associated audio data, in order to identify and display the video stream(s) of the participant(s) who are currently speaking, reacting, or doing.

[0091] In some embodiments, this may be the only consideration when
20 implementing follow the action storytelling. However, in other embodiments, weight can be given to other factors in selecting the video streams to be displayed in the follow-the-action mode. For example, a user may be interested in a non-verbal event associated with one or more of the other users. This can include an event such as a detected change in the number of users that are present in the video image received from a particular user device.
25 It may be desirable to display this kind of event to a user, in addition to, or in place of, the video images associated with verbal-events. That is, call video images for display may be selected based on their visual content, or various activity detection as well as or instead of any voice data associated with those images. In so doing, a user can view the activity of the other users in the video call (be it verbal or non-verbal) in a story-like manner. It may be
30 desirable to display non-verbal events for a limited time interval; for example, to ensure that a user's focus is not taken away from a speaking user for too long.

[0092] As another example of a non-verbal event that can trigger the selection of a video stream is a change in the facial expression of one or more of the other users, which may have occurred as a result of one of the users reacting to a speaking user's speech. It

may be desirable to display one or more of these reactions, as and when they occur, so as to enable a user to view the activity of the other users in the video call, in a story-like manner. Likewise, these reactions may also be displayed for a limited time interval; for example, to ensure that a user's focus is not taken away from a speaking user for too long.

5 **[0093]** Examples of follow-the-action story telling based on non-verbal cues are described in the Applicant's co-pending US patent application numbers 15/048800 and 15/048750, both of which are incorporated herein by reference in their entirety.

[0094] There are additional advantages to follow-the action story telling. As well as providing an improved user experience, follow-the-action storytelling also ensures that
10 the user device does not allocate display resources to video data that does not add to the user's experience of the video call, which is particularly, though not exclusively applicable, to mobile, tablet or certain laptop devices within limited available total display areas. Moreover, in preferred embodiments of the present invention, the video streams carrying the video images are transmitted to a central relay node, for example a relay server or a user
15 device of one of the participants which is hosting the video call. The relay node selects one or more of the streams to be displayed to each participant, and transmits only those stream(s) to that participant (where the relay node is a user device, at least one of the streams eligible for selection can be generated local at that user device). That is, only as many video streams are needed are transmitted to each participant. This saves bandwidth whilst at the same time
20 improving the user experience.

[0095] Notwithstanding, the inventor of the present invention has recognized that, whilst follow-the-action storytelling can provide a highly engaging experience for users in many circumstances – by focusing-in on the current action wherever that may be – there are certain circumstances in which an alternative user experience may in fact be more engaging.
25 In particular, he has recognized that a participant who is late to join a video call (late joiner), i.e. when multiple participants have already joined and have been conducting the conference for some time, may not benefit fully from follow the action storytelling initially.

[0096] The described embodiments of the present invention provide the late joiner at his user device – for an initial duration commencing upon his joining (start phase) – with
30 what is referred to herein as a start storied phase experience, which is unique to that late joiner.

[0097] Detection of a remote participants "late to join" status instigates this start storied phase experience, in which a start phase multi grid introduction to the current participants multi modal experience is provided at his user device, so as to increase the late

joiners awareness of attendees, locations and current activities. Upon expiry of the initial duration, his user devices switches to a follow the action story view phase, bringing his experience into line with that of the existing participants.

5 [0098] In other words, when a late participant first joins a video call, a late-to-join mode (introductory mode) is selected for an initial duration for that participant. In the late-to-join mode, the focus is on providing the late joiner with a full and comprehensive overview of the current state of the call. For example, video images of all (or the majority) of the participants may be displayed to the user, even in-active participants who are not currently part of the action, i.e. who are not currently active speakers and who are not driving
10 the call with non-verbal cues. This improves activity and people engagement/awareness by giving a visual overview of all people and activity currently progressing in the call. In addition, avatars may be displayed in this initial duration representing any participants who are unable to provide video streams, for example because their user devices lack the necessary capability of because they are operating under restricted network conditions, this
15 is an effective means of drawing the late joiner into the call.

[0099] After the initial duration, a follow-the-action mode (focused engagement mode) is selected for the late joiner, so that for a subsequent duration a limited number of video images is selected intelligently and displayed to him, bringing his experience into line with the existing participants. Thus, allowing the late joiner that is now –activity and group
20 aware- to participate in the group activity at an equal engagement follow the action behaviour level of awareness.

[0100] The initial duration (i.e. late-to-join phase) may be several seconds (at least two seconds, and possibly more – e.g. five to ten seconds). The subsequent duration (i.e. the follow-the-action phase) will typically last longer in practice, e.g. at least a minute.

25 [0101] Figure 1 shows a communication system 100 comprising a network 2, to which a relay server 102 and a plurality of user devices 6.1,...,6.M+1 (M+1 in total) are connected. The user devices and the server are nodes of the network 2. Each of the user devices 6.m (where m denoted the mth user device) is operated by a respective user 4.m.

30 [0102] The users 4.1,...,4.M+1 are participants in a video call conducted via the network. User M+1 is a late joiner who has joined the call after users 4.1,...,4.M, who are existing participants of the call at that time.

[0103] Each of the user devices 6.m has at least one image capture device 7.m (camera), for use in capturing a video image of user 4.m that is transmitted from that user device 6.m to the network 2 in a video stream of the video call.

[0104] Each of the user devices 6.m comprises a processor 5.m, such as a CPU or CPUs (e.g. in a multi-core processor), on which an instance of a communication client is executed. The communication client establishes the necessary connections over the network 2 to allow the users to participate in the video call using the user devices.

5 [0105] The communication client on each user devices 6.m generates at least one video stream for transmission via the network 2 as part of the video call. The video stream carries a moving image generated using the at least one camera 7.m of that user device 6.m, by applying video encoding to video data captured by the camera 7.m.

10 [0106] The network 2 is a packet-based communication network, such as the Internet.

[0107] The users 4.1,...4.M+1 are users of the communication system 100. The communication system 100 may have other (possibly many) users who are not shown in figure 1. Whilst in figure 1, each user device is operated by one user, in some cases at least one of the user devices may be operated by multiple users. That is, multiple users may participate in the video call using the same user device, which can for example be a video conferencing device. Moreover, not every user device in the call needs to have a camera – that is, there may be additional users operating additional user devices and participating in the video call even if their user devices lack a camera, or lack the capability to provide a call video stream for some other reason, for example because they have limited available network bandwidth.

20 [0108] Moreover, whilst each user device 6.m is shown as having one image capture device 7.m, one or more of the user devices 6.m may have one or more additional camera devices allowing multiple video streams to be generated and transmitted from that device. The additional camera device can for example provide an alternative angle or view (e.g. wide-view) from which to capture a video image of the user(s). Moreover, at least one of the user devices may generate multiple video streams using the same camera. For example, a single camera may generate multiple video streams, each derived from a different region of a high-resolution master video image captured with a single camera.

30 [0109] In any event, N video streams are generated by the M user devices having video-stream capability. In the simplest case, $N=M$ (one stream per user device with video capabilities), but as noted in other cases $N>M$ (multiple streams generated by at least one user device, from the same or different cameras). The N video streams are denoted s_1, \dots, s_N below. In this example there are M participants (one per device), however in other cases there can be more participants than user devices (when multiple users are using at least one

of the devices), or even fewer participant that user devices (where one participant is using multiple devices). M may be the total number of user devices, or there may be at least one additional user device in the call which lacks such video stream capability.

5 [0110] In the following examples, the number of streams N is at least four and may be considerably higher for a video call with many participants.

[0111] In the described examples, the video streams are transmitted to a relay node, which can be the relay server 102. Alternatively, it can be one of the user devices 6.1,...,6.M, e.g. the user device of the participant who is hosting the video call, which in some implementations is automatically the participant who instigated the video call.

10 [0112] Figure 2 shows a functional block diagram of a media relay system 200, which represents functionality implemented by the really node.

[0113] Where the relay node is the server 102, this functionality can be implemented by server code executed on at least one processor 105 of the relay server 120. That is, server code executed on one processor, or whose execution is distributed across multiple
15 processors. The multiple processors can be part of one computer device, or multiple computer devices which may or may not be spatially collocated. For example, the relay server 102 may be implemented in a cloud computing environment, wherein its functionality is implemented by one on more virtual devices, which in turn run on one or more physical computer devices which can be in the same or different data centres where there are multiple
20 computer devices.

[0114] Where the relay node is one of the user devices themselves, the functionality of the media relay system 300 is part of the functionality of the communication client executed on the processor of that user device.

25 [0115] Each time a new participant joins the call, a notification of the new participant is received at the relay node. This can for example be triggered by the detection of an initial receipt of at least one video stream from the new participant at the relay node.

[0116] The media relay system 300 receives the N video streams s_1, \dots, s_N from the M devices 6.1,...,6.M. When the relay node is the server 102, these are all received via the network 2 from the user devices 6.1,...,6.M. When the relay node is one of the user devices,
30 at least one of those streams may be generated locally at that user device and the rest received via the network 2 from the remaining user devices.

[0117] The system 200 is shown to comprise functional modules representing different parts of the system's functionality, namely a stream selector 202 and a mode selector 204 which operates as a controller of the stream selector 202. Blocks #1 to # M (M

in total) of the stream selector 202 denote individual stream selection functionality for the user devices 6.1,...6.M. That is, a set of one or more of the streams s_1, \dots, s_N is selected for each of the user devices individually, where $S.m$ denotes the set of stream(s) selected for device 6.m operated by user 4.m. This allows each participant to be provided with a unique
5 experience tailored to him.

[0118] Mode selector 204 selects a mode for each of the user devices 6.m depending on when its user 4.m joined the video call. In this example, the system 300 is shown at a time when the follow-the-action mode is selected for each of the existing participants 4.1,...4.M, whereas the late-to-join mode is selected for the late joiner 4.M+1 who has
10 recently joined the call.

[0119] The late joiner 4.M+1 can be identified as such by mode selector 204 based on a time at which he joins the communication event. For example, based on an elapsed time of the communication event that the join time (i.e. how long it has already been going on for), the number of exiting participants M at that time (i.e. how many people have joined
15 before him), some other late-to-join condition, or any combination thereof.

[0120] User device 6.m is configured to display a user interface (500 – see below), which is controlled to operate in the mode selected for that device 6.m.

[0121] For each of the existing participants 4.1,..4.M, the selection of the follow-the-action mode causes between one or two streams to be selected for displaying to that
20 user, by blocks #1,...,#M respectively, in accordance with the follow-the action principles set up above. In addition, display data $D.1, \dots, D.M$ is generated by that block a layout for displaying the video image(s) carried by those selected stream(s) at the relevant user device, in the manner described below.

[0122] For the late joiner 4.M+1, the late-to-join mode remains selected for an initial
25 duration after his joining. The selection of the late-to-join mode causes a greater number of streams (between three and four) to be selected for that user at a given time whilst that mode remains selected. Display data $D.M+1$ defining a layout for displaying the selected streams $S.M+1$ is also generated.

[0123] Each participant's own video stream(s) is not transmitted back to him over
30 the network 2, thus the total number of video streams eligible for selection for the late joiner 4.M+1 is N less the number of the late joiner's own video streams (if any – as noted, not every participant need to have video capture capabilities). For example, if one of the streams is the late joiner's, then there are N-1 eligible for selection for displaying to the late joiner.

[0124] Although not eligible for selection in this manner, a self-video image of the late joiner's own call video may be locally displayed at his device, for example in a corner of the display area in a conventional manner. The same goes for the other participants. The “me” video (i.e. self-image) of the user’s own video is displayed in the grid when “pinned”
5 by the user or constrained to the grid on mobile devices.

[0125] For each participant 4.m, the set of stream(s) S.m selected for that participant is transmitted to that participant's user device 6.m along with the display data D.m via the network 2. An exception to this is when the relay node is one of the user devices, in which case the streams selected for that user device are outputted locally instead without needing
10 to be transmitted via the network 2.

[0126] The selected stream(s) S.m are displayed at user device 6.m according to one of four pre-determined layouts 402 (one stream), 404 (two streams), 406 (three streams) and 408 (four streams), which are illustrated in figure 3.

[0127] Layout 402 has only one display region for displaying only one call video
15 image at a time, which occupies all of a call video display area of a display of user device 6.m. The call video area can be the whole of an available display area of the display, or a portion of the available display area currently reserved for call video (or more generally video of the communication event).

[0128] Layout 404 has two display regions for displaying two call video images at
20 a time (i.e. simultaneously) such that each occupies half of the call video display area. These are displayed side-by-side in this example, but could alternatively be displayed one-above-the other.

[0129] Layout 406 has three display regions for displaying three call video images
25 at a time such that one of call video images occupies a greater portion of the call video display area than each of the other two images. That is, one of the call video images is displayed in larger display area than the other two. Each of the other two images occupies half of the remaining call video display area.

[0130] Layout 408 has four display regions for displaying four call video images at
30 a time in a two-by-two grid arrangement, with each occupying a quarter of the call video display area.

[0131] Layouts 406 and 408 (three and four streams) are used only for the late-to-join mode, whereas layouts 402 and 404 (one and two streams) are used only for the follow-the-action mode.

[0132] Figures 5A-5D shows a portion of graphical user interface 500 of the client application, which is rendered in a display area 502 of the display of user device 6.m, in order to display the selected stream(s) S.m according to the received display data D.m within the display area 502. In each of figures 5A-5D, the user interface 500 is shown operating
5 in a different one of four call video layout modes, in which to layouts 402, 404, 406 and 408 respectively are rendered by the user interface.

[0133] Denoting the number of call video images displayed simultaneously at time t as $Q(t)$, it can be seen that:

[0134] $Q(t)=1$ at any time t at which layout 402 is selected (as in figure 4D);

10 [0135] $Q(t)=2$ at any time t at which layout 404 is selected (as in figure 4C);

[0136] $Q(t)=3$ at any time t at which layout 406 is selected (as in figure 4B);

[0137] $Q(t)=4$ at any time t at which layout 408 is selected (as in figure 4A).

[0138] For example, when operating in a persistent stage view mode, content in "stage" plus four active video positions and me (i.e. the participant's self-image) is a core
15 view we design for as well. That is, (up to) four call video images in the call video display area, a self-video image and a (shared) content video image (outside of the call video display area) are displayed simultaneously – an example is shown in figures 7A and 7B, and describe in further detail below.

[0139] When operating in a people/content toggle mode, up to four grid videos can
20 be shown in stage while live content can be shown in bottom row with me and passive participants. Thus a passive way to show content as a participant as well.

[0140] As noted above, embodiments of the present invention can provide what is referred to as a content-as-a-participant experience, wherein one or more active content views can populate one or more of the available grid views depending on user or system
25 designated importance. For example, if current activity is detected in a shared document or if a shared desktop is active in the grid views a late to join curated view can populate such activity like a person video stream.

[0141] By selecting, for the late joiner 4.M+1, only layouts 406 and 408 for the initial duration in which the late-to-join mode is selected, a minimum and a maximum
30 number of call video images is thereby selected for that duration, which are three and four respectively. Because only those layouts are used for the initial duration, the total number of simultaneously displayed call video images remains between three and four inclusive at all times during the initial duration. However, in other implementations this can be relaxed, and all four layouts may be eligible for selection in the late-to-join mode.

[0142] Likewise, by selecting for the late joiner 4.M+1 only layouts 402 and 404 for the subsequent duration, a minimum and maximum number of call video images is selected for that duration, which are one and two respectively. Because only those layouts are used for the subsequent duration, the total number of simultaneously displayed call video images
5 remains between one and two inclusive at all times during the initial duration. The same goes for the existing participants 4.1,...,4.M when the follow-the-action mode is selected for them.

[0143] The call video layout mode for participant 4.m is set by display data D.m, which can for example specify the selected layout and may also specify which of the video
10 streams in the set S.m should be displayed in which regions of that layout.

[0144] Over the course of initial duration, more than four of the N streams plus content (where applicable, e.g. in the persistent stage view) may be displayed in total (but no more than four at any one time) according to the three and four video layouts 406 and 408. For example, during a first part of the initial duration the four-video layout may be
15 selected by stream selector block #M+1 for the late joiner, and cycle through the various video streams of the other participants. Once all or most of these have been displayed to the user, model selector block #M+1 may switch to the three-video layout for the final part of the initial duration to provide a seamless transition to the follow-the action mode. Upon transitioning to the follow-the action mode, stream selector block #M+1 transitions to two
20 or one video streams, displayed according to the two and one video layouts 404 and 402 respectively.

[0145] Any participants who are not providing video streams, may still be represented in the late-to-join mode, for example as avatars displayed simultaneously with the video streams. In some cases, these avatars are only displayed for the late-to-join mode,
25 to provide an initial overview of all participant including those without video, but are not displayed in the follow-the-action mode as they could be distracting in the long run.

[0146] Note also that, in some cases, at least some of the video streams may not be generated using cameras. For example, at least one of the video streams may implement another video modality, such as screen sharing, document sharing or a shared whiteboard
30 or note pad. In the late-to-join mode, these can be displayed in the same manner as video streams. In the follow-the-action mode, this may only be displayed if it is currently part of the action, for example the two-video layout may be used to shown a shared screen or document simultaneously with the video of one of the participants who is giving a presentation on it. To implement this, each video stream may also carry a modality

indicator, which is used to select the stream(s) S.m for displaying in dependence on the current mode for block #m of stream selector 202.

[0147] The displaying of the greater number of video streams during the initial duration may require more bandwidth, as more video streams have to be transmitted from the relay node to the new participant's user device 6.M+1 (as compared with follow-the-action story telling). However, this additional bandwidth is only needed for as long as it takes to immerse the new participant 4.N+1 in the video call, at which point bandwidth can be saved by switching to the follow-the action mode, thereby reducing the number of video streams transmitted to him without impacting on his user experience. The switch takes place automatically, in response to expiry of the initial duration. This can be a predetermined duration or determined dynamically in dependence on the communication event. In any event, the switch takes place automatically based on a duration timer.

[0148] As shown in figure 5, for the existing participants 4.1,...4.N in the follow-the-action mode, a visual notification 504 of the new joiner is briefly displayed on the user interface 500. This keeps them informed of new joiners, without being overly distracting.

[0149] The visual notification 504 can for example comprise an avatar or image of the new joiner captured with his camera.

[0150] Figure 6A shows a first example of the client GUI in greater detail. As can be seen in figure 6A, the display area 502 in which the selected video stream(s) are displayed is only a portion of the display area. Outside of this portion 502, another portion of the client GUI is displayed, which comprises a self-video image 506 of the user 4.m of device 6.m, a shared content video-image 508 (which the user 4.m may for example be sharing himself), avatars 501 representing other call participants. Function buttons 512 for the call are also displayed within the display area 502 in this example. That is, GUI elements that are selectable to instigate a particular function relating to the communication event, such as terminating the call, muting or unmuting a microphone of user device 6.m, enabling or disabling an image capture device of user device 6.m (so as to instigate or terminate the transmission of a video stream to the other participants), and/or enabling or displaying screen sharing from user device 6.m.

[0151] Note that, in addition or as an alternative to the avatars 501 outside of the display area 502, at least one avatar may be displayed within region 502 for at least part of the late to join phase, instead of a video image derived from one of the video streams. In this case, the total number of images (i.e. avatar plus video images) within the display area 502 is dictated by the selected layout.

[0152] Figure 6B shows a second example of a client GUI which is the same as that of figure 6A, other than the fact that display area 502 in which the selected video stream(s) are displayed is non-rectangular, in order to accommodate the self-video image 506 and shared content 508 at one corner of the GUI. The avatars 510 are not displayed in this
5 example.

[0153] Figures 7A and 7B show an alternative "staged content" configuration of the client GUI 500 when operating in the late-to-join mode and follow-the-action mode respectively. In this example, the shared content 508 takes "centre stage" i.e. is displayed is a large portion of the total available display area, which may be appropriate where the
10 content is the focus of the conference. The display area 502 is a smaller portion of the total available display area, below the shared content. In the late to join mode (figure 7A), four call video images are shown displayed in the display area 502 along with avatars. The function buttons 512 are displayed outside of the display area 502, within the larger portion of the total display area in which the shared content 508 is displayed, such that they are
15 overlaid on the displayed content 508. The self-video image 506 is also displayed outside of the display area 502, towards one corner of the GUI 500. By contrast, in the follow-the-action mode (figure 7B), only one call video image is shown displayed within the same display area 502, and the avatars are not displayed. This frees-up space within the display area 502, which is exploited by moving the function buttons 512 such that they are now
20 displayed within the display area 502. The self-image 506 and the content 508 are displayed at the same locations on the display.

[0154] Sub areas (sections) 502A-B and 502C-D of the display area 502 are marked in figures 7A and 7B respectively, and visual information of the GUI 500 is displayed within these sub-areas as set out below.

25 FIG 7A:

[0155] 502C: Dominant activity video position 1-4 (i.e. four video position within 502C, each occupying an equal portion thereof) which populate from right to left for dominant activity or dominant speaking. Once populated passive participants in section 502D (see below) can be promoted into the active stage positions of 502C via activity
30 duration timer detection, by swapping that participant with the least active stage position participant currently in 502C (who in turn is demoted to passive section 502D). That is, display a call video image of the promoted participant in 502C in place of the demoted participant's, and displaying an avatar of the demoted participant in 502D. Participants can be video or audio only avatar representations.

[0156] 502D: Passive activity grid: an area for passive participant avatar representation to give awareness of meeting size or show pre-promotion activity representation (active speaking prior to promotion to dominant activity view. Populates from right to left and with overflow in large meeting experiences. Representations of transition activity are shown for: join/leave, stand/sit. Speaking, Expressions, attribute states, etc.

FIG 7B:

[0157] 502A: Dominant activity video position: one video at a time is promoted at a time for follow the action dominant activity and/or dominant speaking. Participants can be video or audio only avatar representations.

[0158] 502B: Passive activity grid. Area for passive participant avatar representation to give awareness of meeting size or show pre-promotion activity representation (active speaking prior to promotion to dominant activity view. Populates from right to left and with overflow in large meeting experiences. Representations of transition activity are shown for: join/leave, stand/sit. Speaking, Expressions, attribute states, etc. If passive participants occupy grid zone 502B, call controls 512 bump vertically over lower content grid during user control on state. In other words, although not shown in figure 7B, **avatar** images can be displayed in area 502B; however, the number of **video** images in the area 502 is limited to (in this case) one.

[0159] Note that each avatar image displayed in passive regions 502B/502D occupies a smaller portion of the display area than any avatar image displayed in active regions 502A/502C.

[0160] The shared content 508 can for example be a screen share video image, or shared document video image.

[0161] A consequence of the above regime is that fewer call video images are displayed by the new participant's user device $6.M+1$ at any one time during the subsequent duration than at any one time during the initial duration. In other words, adopting the $Q(t)$ notation defined above:

$$Q(t_1) \geq Q(t_2) + 1$$

[0162] for every single time t_1 in the initial duration and every single time t_2 in the subsequent duration, where " \geq " denotes strictly greater than or equal to. This holds true even when the video call has at least as many participants, whose user devices do have video capture capability, throughout the subsequent duration as for the initial duration. That is –

denoting the number of participants in the communication event whose devices have video stream capability at time t as $P(t)$ – even when:

$$P(t_1) \leq P(t_2)$$

5 [0163] for every single time t_1 in the initial duration and every single time t_2 in the subsequent duration. In other words, the number of call video images displayed to the late joiner is not decreasing because the number of participant's operating user devices with video-stream generation capabilities is decreasing, but is decreasing as a deliberate choice independently of this. For example, all of the existing participants who are participants when the late joiner joins may remain participants throughout the initial and subsequent
10 durations.

[0164] However, as noted in other implementations, the restriction on the late-to-join mode may be relaxed, allowing all four layouts to be selected – and hence between one and four images/avatars to be displayed in the grid view – in the late to join mode.

15 [0165] For example, in the late-to-join mode, there may be a default grid view that would be utilized a large percentage of the time, and all stage grid views 1-4 may be available for the late to join experience, that are populated by people, groups or content. The selection logic (mode selector 204) is not constrained to just one of these grids only in this mode.

20 [0166] When a passive avatar row is created and populated from right to left in the bottom row next to the live stage videos, dominant speaking video, or me video to the left, an “overflow” position may be created in the furthest left position in the event of large meetings. This allows for large meetings with more participants than positions available to be represented as +2, +3, +4 and so on in the lower far left passive position... As the conversation unfolds, the most active participants populate in the active and passive
25 individual representation positions via activity points priority stacking. Active and passive positions can be populated by individual, group representation or video.

[0167] Note that whilst in the above the streams are transmitted to a relay node for selection, this is not essential. Each user device can for example transmit its stream(s) directly to every other user device, and each user device may perform its own selection of a
30 subset of the received streams to be displayed. Although less preferred in some contexts due to the potentially higher bandwidth requirements (due to transmitting streams that are not going to be displayed), it may still be viable in other contexts, and has the advantage that no relay node is needed. The burden imposed on the relay node can be high in terms of

required processing and bandwidth resources, particularly where it is receiving a large number of streams.

5 **[0168]** As another example, user device 6.m can control which streams it receives based on stream subscription. That is, by subscribing to desired streams from the relay node. For example, the user device 6.m may subscribe to the video stream of the actively speaking participant when its user interface is operating in the active-speaker mode, by transmitting at least one active speaker subscription message to a relay node. In some cases, it may only subscribe to that video stream in that mode, at least some of the time. The user device does not need to know who the active speaker is (that can be determined at the relay
10 node); it can simply send a stream subscription request requesting the stream of the active speaker whoever that may be. If the active speaker changes the user device will automatically receive the video stream of the new active speaker without having to alter its subscription.

15 **[0169]** In the late-to-join mode, the user device may subscribe to all (or some) of the video streams, by transmitting at least one stream subscription message to the relay node, wherein the client receives the subscribed-to video streams from the relay node via the network. More generally, it may send a stream subscription request requesting the video stream of at least one of the participants who is not an active speaker, and receive the subscribed-to video stream at a time when that participant is not the active speaker.

20 **[0170]** Whilst the above has been described with reference to specific examples, other variations are within the spirit and scope of the present invention. The scope is not in limited by the examples described above, but only by the following claims.

Claims

1. A method of controlling a user interface of a user device operated by a new participant of an established communication event, the established communication event being conducted between a plurality of existing participants by transmitting and receiving video streams via a communication network between user devices operated by those existing participants, the method comprising steps of:

receiving a notification of the new participant joining the established communication event;

determining, based on the received notification, whether the new participant satisfies a late-to-join condition; and

if the late-to-join condition is determined to be satisfied: controlling the user interface to operate in a late-to-join mode for an initial duration, and in response to expiry of the initial duration, controlling the user interface to operate in an active speaker mode instead;

wherein controlling the user interface to operate in the active speaker mode comprises selecting for the new participant, based on audio data generated at the user device of an actively speaking one of the participants: one of the video streams of the actively speaking participant, or an avatar image representing the actively speaking participant, thereby causing the user interface to display the selected avatar image or a video image derived from the selected stream.

2. A method according to claim 1, wherein controlling the user interface to operate in the late-to-join mode comprises selecting for the new participant one or more of said video streams and/or one or more avatar images independently of any audio data generated at the user devices, thereby causing the user interface to display the selected avatar images or one or more video images derived from the selected video streams for at least part of the initial duration.

3. A method according to claim 1 or 2, wherein no more than a maximum number of video images derived from the video streams is displayed by the user interface at any one time when operating in the active speaker mode, and more than that maximum number of video images are displayed simultaneously by the user interface operating in the late-to-join mode at least once during the initial duration.

4. A method according to any preceding claim, wherein controlling the user interface to operate in the late-to-join mode comprises controlling the user interface to

display over the course of the initial duration, for every one of the existing participants, a respective video image derived from one of the video streams or a respective avatar image.

5. A method according to any preceding claim, wherein the late-to-join condition is based on at least one of an elapsed time of the communication event at a join time of the new participant, or a number of existing participants of the communication event at the join time.

6. A method according to any preceding claim, further comprising a step of controlling the user interface to operate in the active speaker mode immediately if the late-to-join condition is determined not to be satisfied.

7. A method according to any preceding claim, wherein the steps are implemented at a relay node other than the new participant's user device.

8. A method according to claim 7, wherein the controlling step is implemented by transmitting from the relay node via the network to the new participant's user device any of the video streams needed by the user interface operating in the active speaker mode to display the video image.

9. A method according to any preceding claim, wherein the relay node is a relay server, or the user device operated by one of the existing participants.

10. A method according to any preceding claim, wherein the steps are implemented by the new participant's user device.

11. A method according to any preceding claim, wherein the user device subscribes to the video stream of the actively speaking participant when the user interface is operating in the active-speaker mode, wherein the user device receives the subscribed-to stream from a relay node via the network.

12. A method according to any preceding claim, wherein in the late-to-join mode, the user device subscribes to at least one of the video streams of a non-speaking participant wherein the user device receives the subscribed-to video stream from the relay node via the network.

13. A method according to any preceding claim, wherein the video image is:
a call video image, derived from one of the video streams, of the actively speaking participant and is selected for the subsequent duration based on the active speaker's audio data.

14. A method according to any preceding claim, wherein at least one video image is displayed by the user interface operating in the late-to-join mode during the initial

duration, the at least one video image being a call video image of a participant who is not actively speaking at that time.

15. A method according to any preceding claim, wherein different video or avatar images are displayed at different times during the initial duration by the user interface operating in the late-to-join mode.

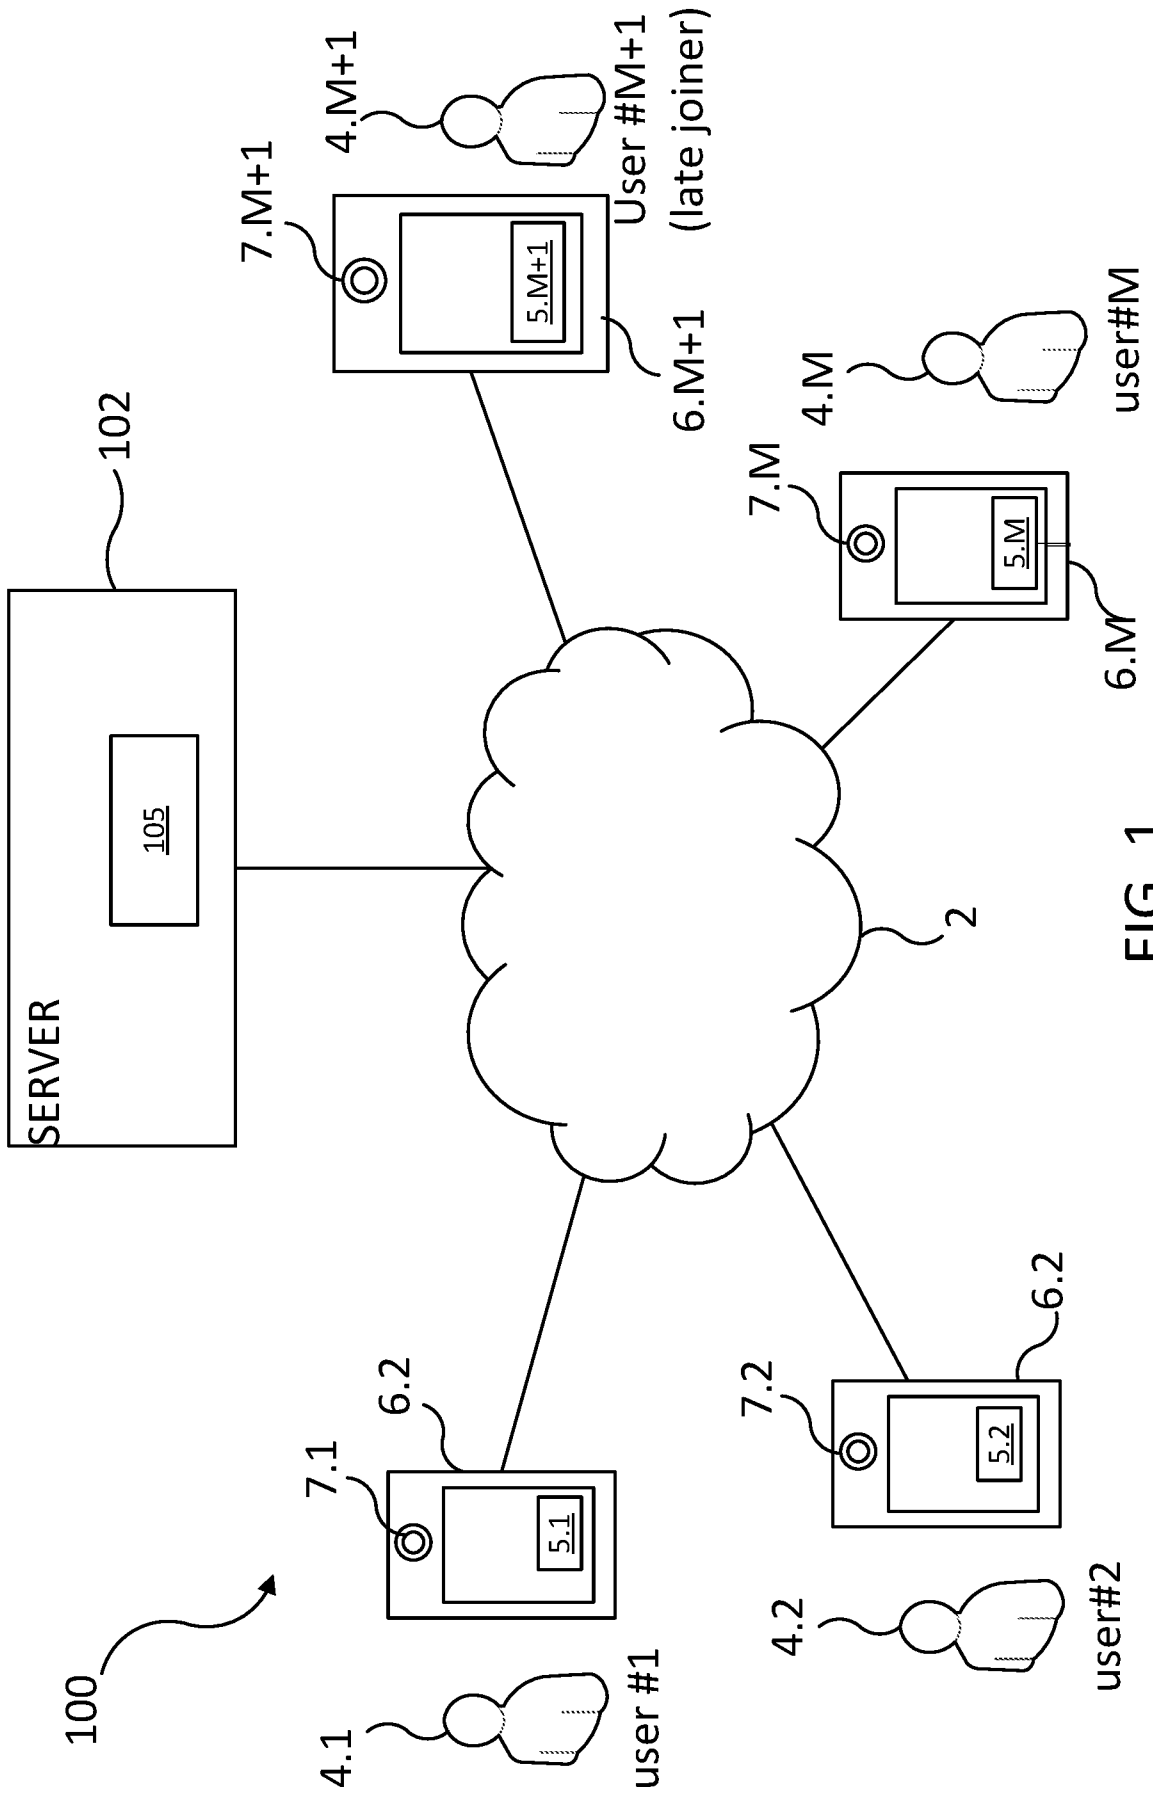


FIG. 1

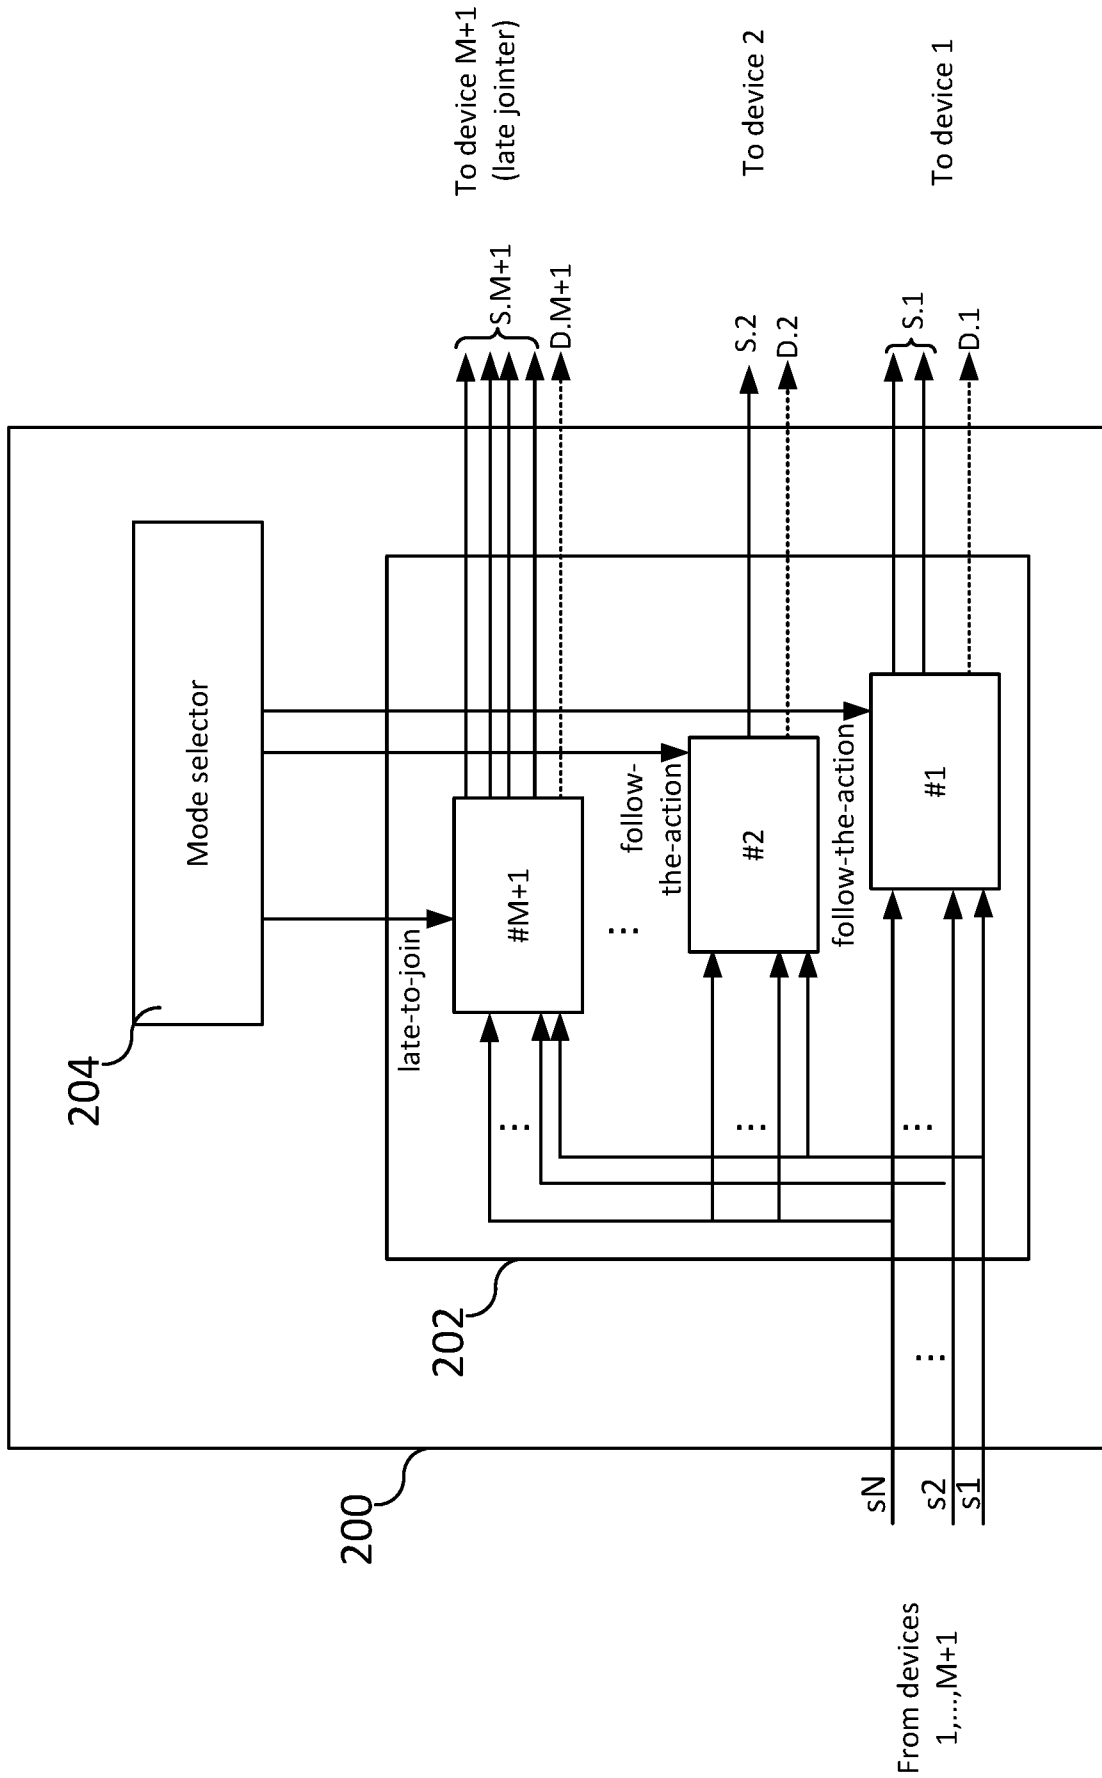


FIG. 2

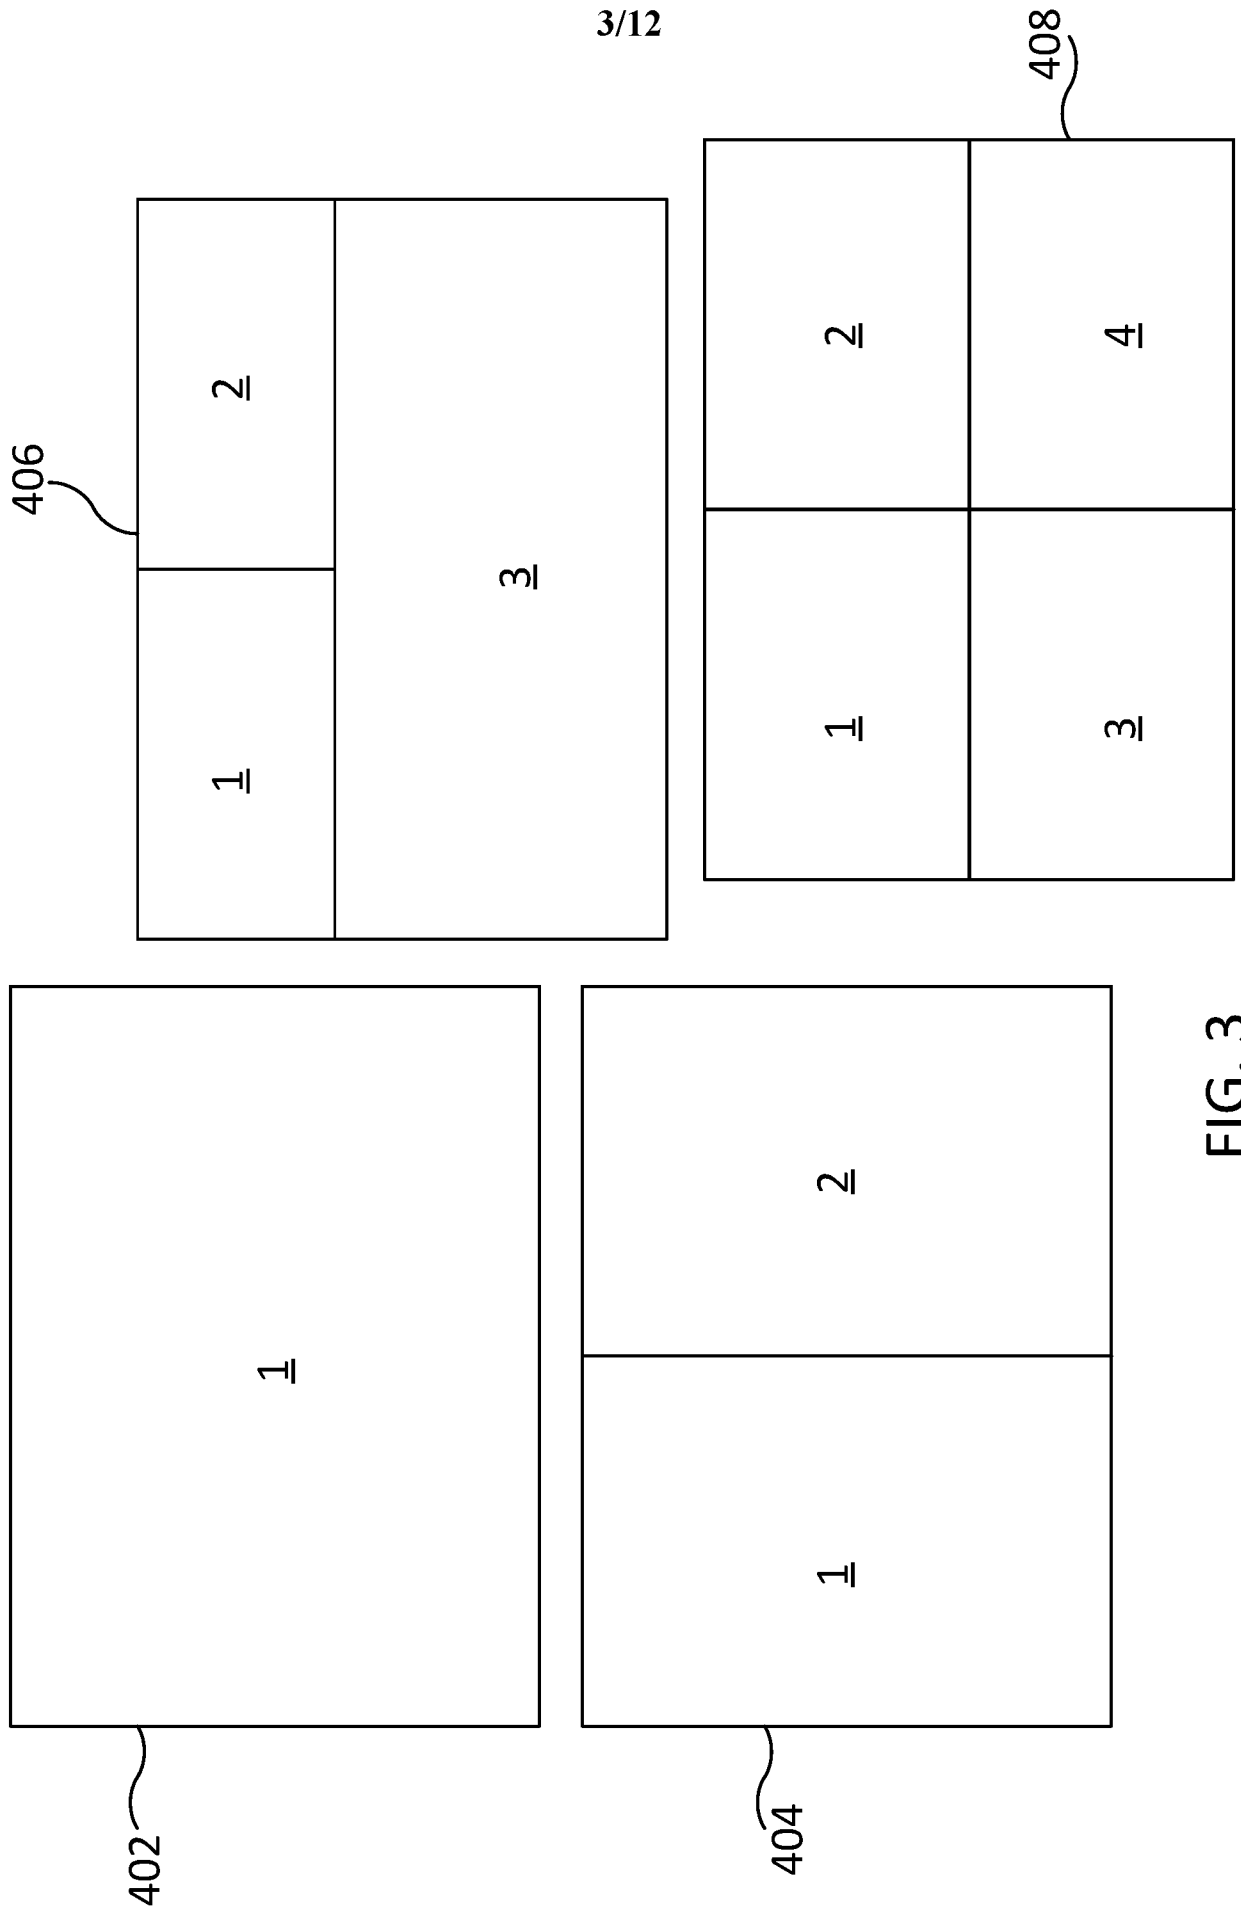


FIG. 3




FIG. 4A




FIG. 4B

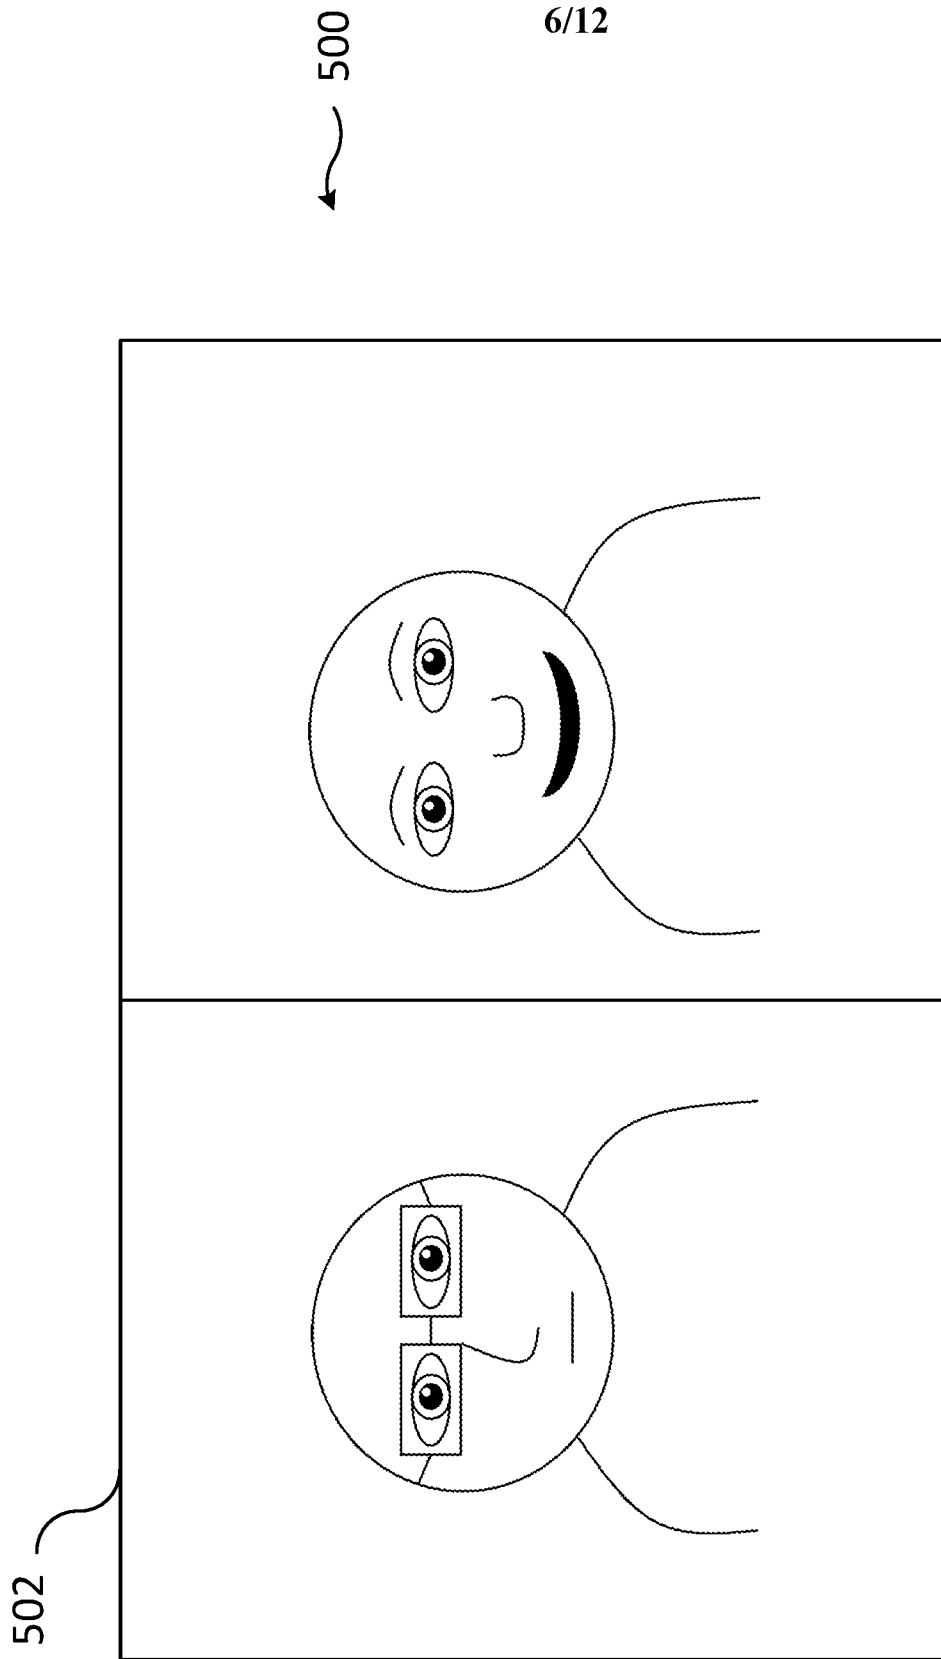


FIG. 4C

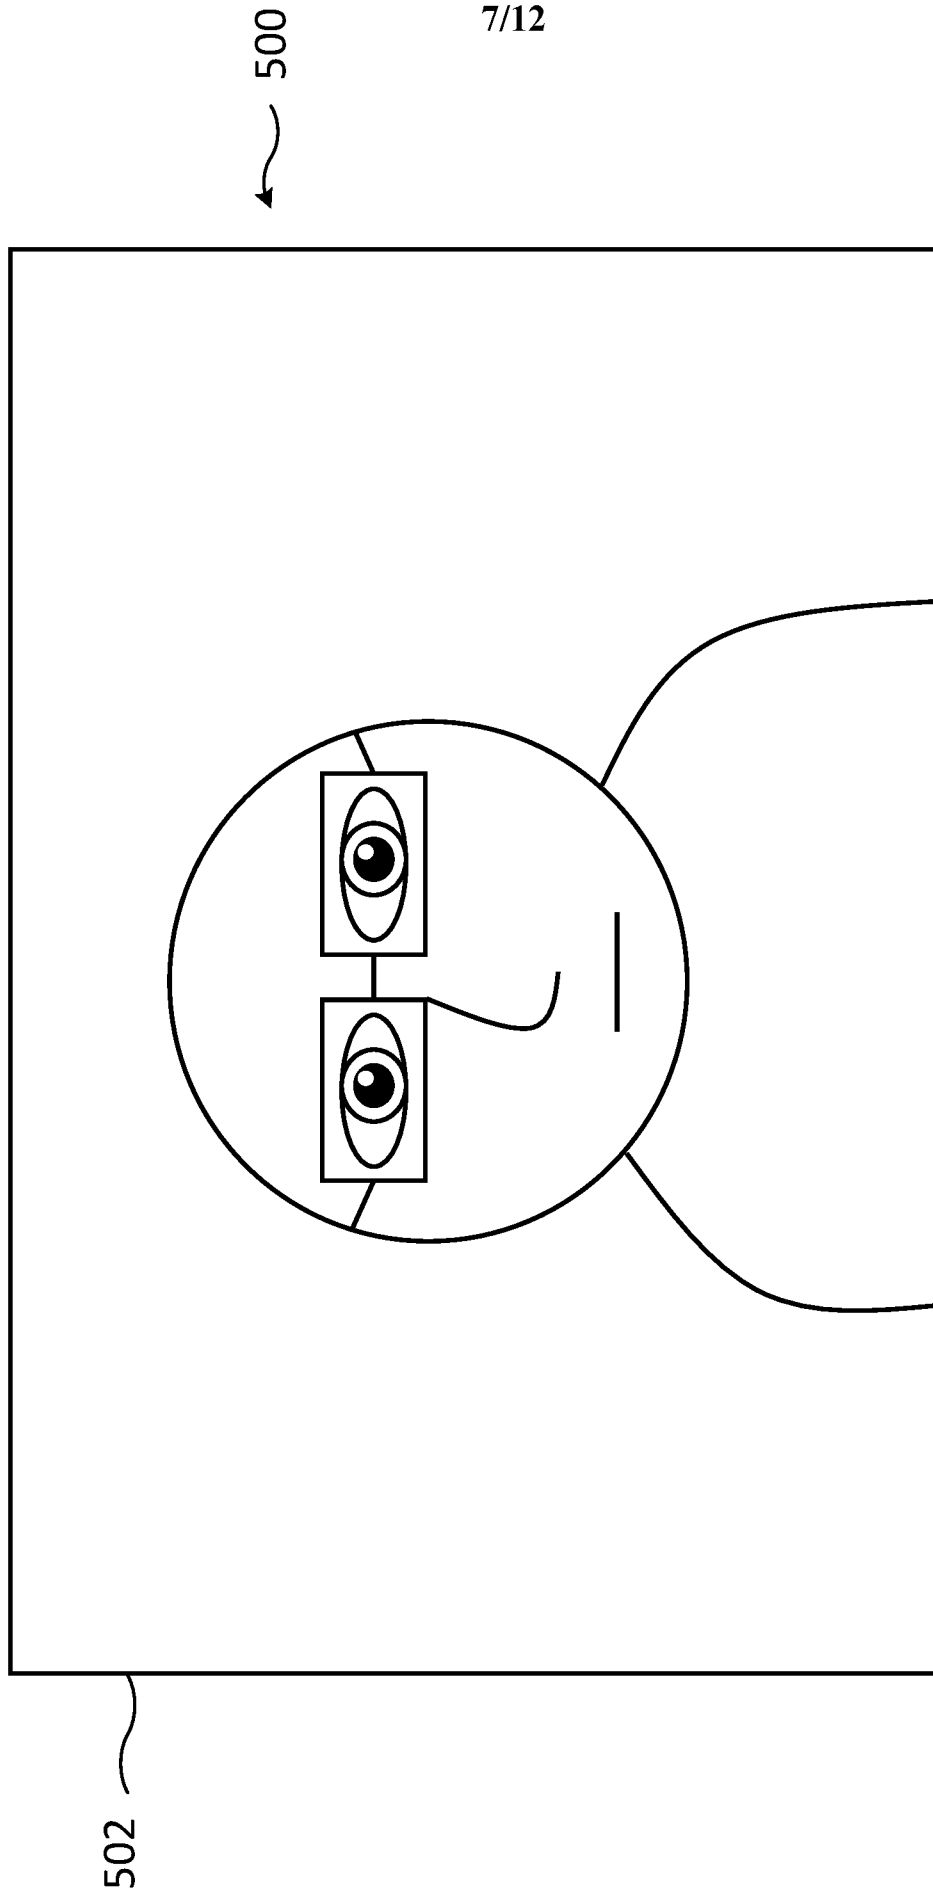


FIG. 4D

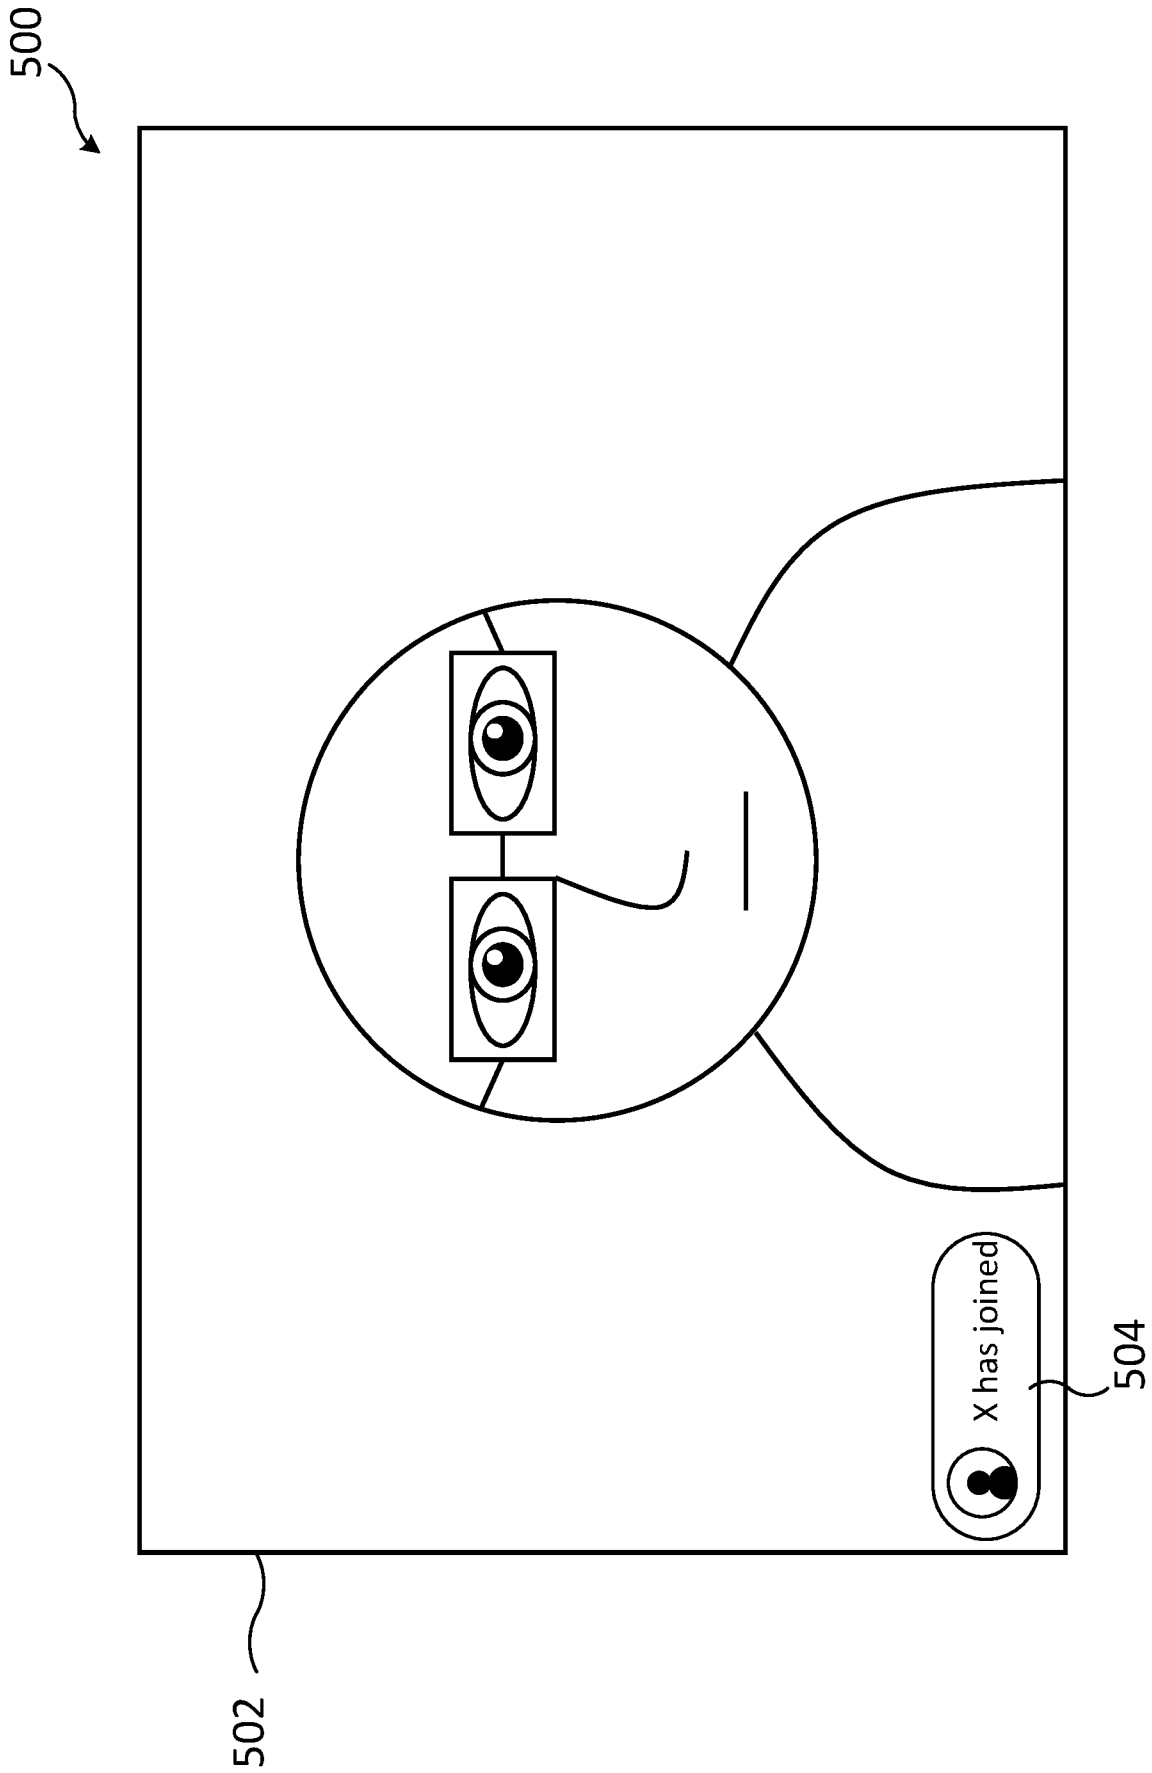


FIG. 5

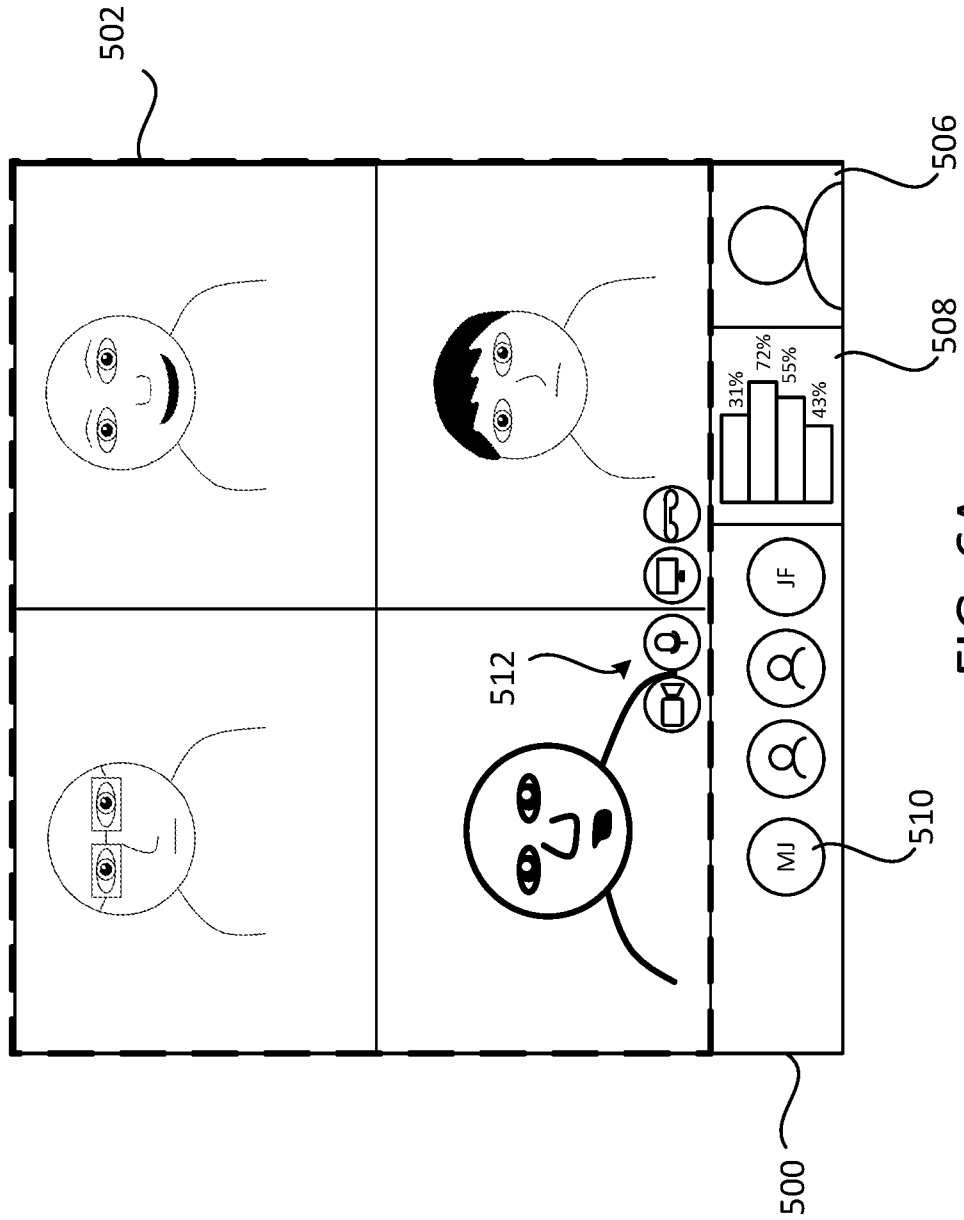


FIG. 6A

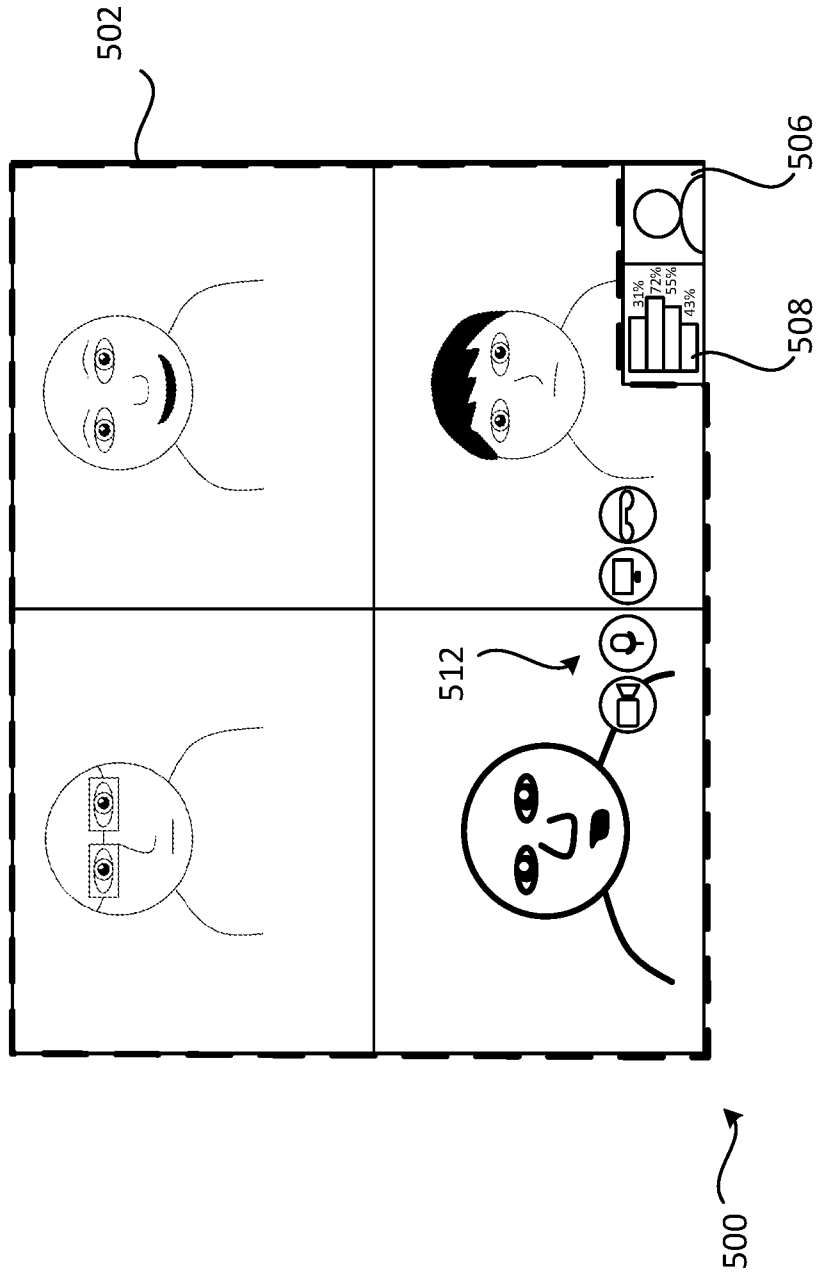


FIG. 6B

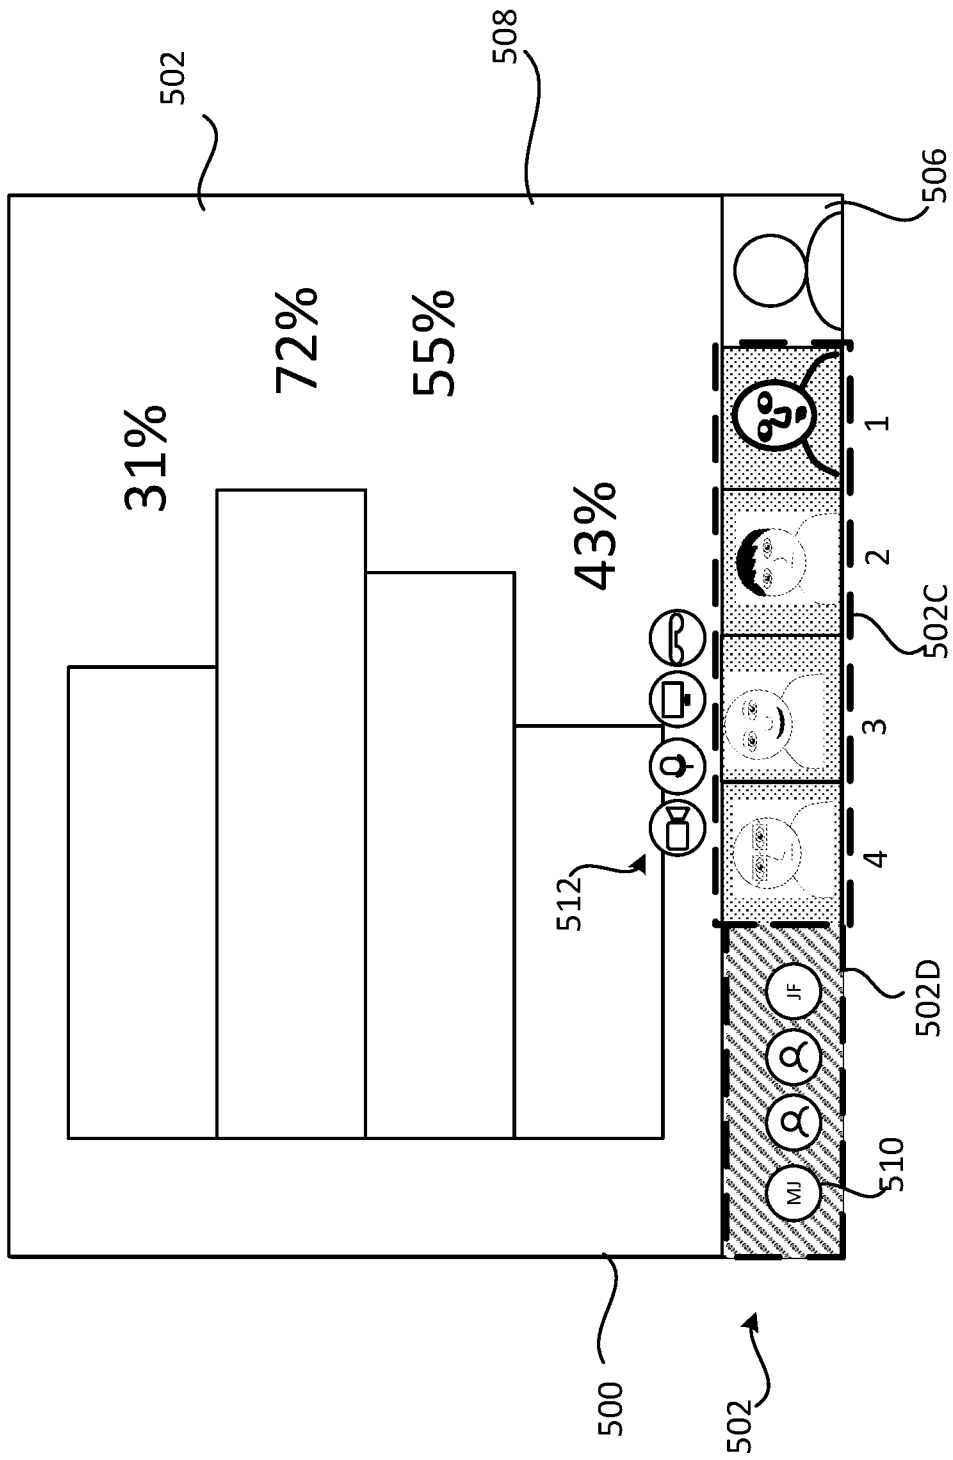


FIG. 7A

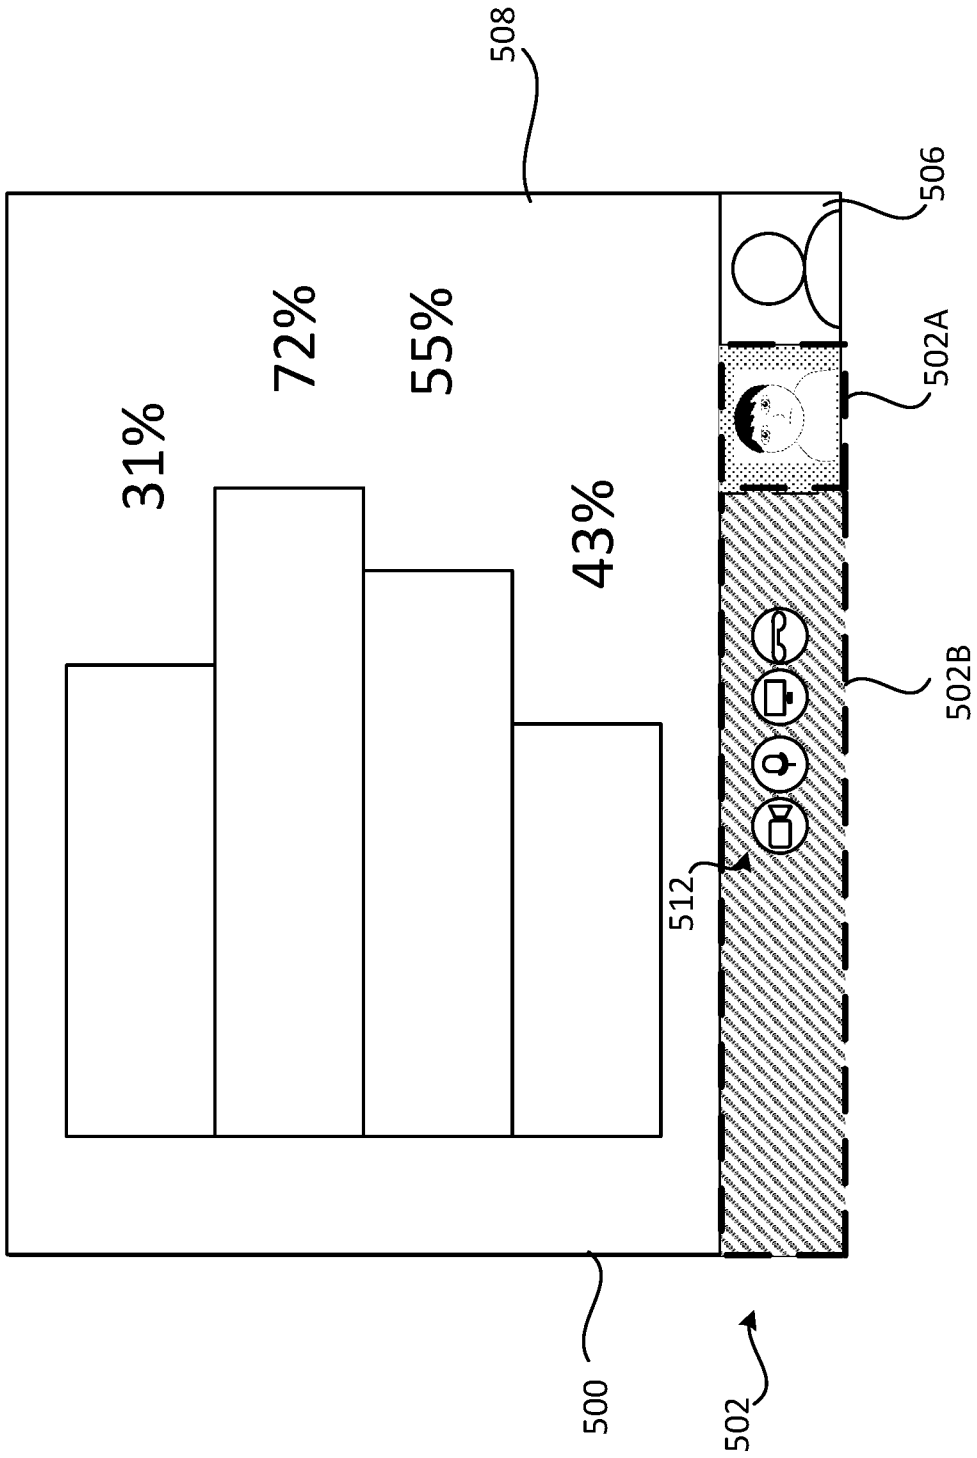


FIG. 7B

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2017/044062

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04N7/15 H04L29/06
ADD.
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED
Minimum documentation searched (classification system followed by classification symbols)
H04N H04L
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 7 308 476 B2 (IBM [US]) 11 December 2007 (2007-12-11) abstract; claims 1-38; figures 1-6 column 1, line 7 - column 2, line 19 column 2, line 50 - column 8, line 13 -----	1-15
X	US 2010/315484 A1 (RAMANATHAN RAJESH [US] ET AL) 16 December 2010 (2010-12-16) abstract; claims 1-20; figures 1-12 paragraphs: 0001-0003, 0019-0059, 0065-0073 -----	1-15

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"E" earlier application or patent but published on or after the international filing date	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"O" document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search 28 September 2017	Date of mailing of the international search report 05/10/2017
---	---

Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer Florou, Rafaela
--	--

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2017/044062

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 7308476	B2	11-12-2007	NONE

US 2010315484	A1	16-12-2010	NONE
