PROVIDING REAL-TIME TRANSLATION DURING VIRTUAL CONFERENCES

FIELD

[0001] The present application generally relates to virtual conferences and more specifically relates to providing real-time translation during virtual conferences.

BRIEF DESCRIPTION OF THE DRAWINGS

[0002] The accompanying drawings, which are incorporated into and constitute a part of this specification, illustrate one or more certain examples and, together with the description of the example, serve to explain the principles and implementations of the certain examples.

[0003] Figures 1-4 show example systems for providing real-time translation during virtual conferences;

[0004] Figures 5A-5B show an example graphical user interface ("GUI") for providing real-time translation during virtual conferences;

[0005] Figures 6-7 show example methods for providing real-time translation during virtual conferences; and

[0006] Figure 8 shows an example computing device suitable for use with example systems and methods for providing real-time translation during virtual conferences.

DETAILED DESCRIPTION

[0007] Examples are described herein in the context of providing real-time translation during virtual conferences. Those of ordinary skill in the art will realize that the following description is illustrative only and is not intended to be in any way limiting. Reference will now be made in detail to implementations of examples as illustrated in the accompanying drawings. The same reference indicators will be used throughout the drawings and the following description to refer to the same or like items.

[0008] In the interest of clarity, not all of the routine features of the examples described herein are shown and described. It will, of course, be appreciated that in the development of any such actual implementation, numerous implementation-

specific decisions must be made in order to achieve the developer's specific goals, such as compliance with application- and business-related constraints, and that these specific goals will vary from one implementation to another and from one developer to another.

[0009] During a virtual conference, participants may engage with each other to discuss any matters of interest. Typically such participants will interact in a virtual conference using a camera and microphone, which provides video and audio streams (each a "media" stream) that can be delivered to the other participants by the virtual conference provider and be displayed via the various client devices' displays or speakers. However, because virtual conferences allow participants to connect from anywhere in the world, oftentimes different participants may have different native languages. And while some people are able to converse in multiple languages, others cannot. Thus, in settings where multiple people do not speak a common language, one or more interpreters may be needed to translate for the various participants. However, this is typically costly and inconvenient.

[0010] To enable different participants who speak different languages to converse with each other, a virtual conference provider may provide real-time translation of a speaker from the speaker's native language into the native language of another participant in the virtual conference. The translation may then be provided to the other participant as text presented in the other participant's native language.

[0011] A difficulty with providing such translation services is that a virtual conference provider may host many hundreds or thousands of simultaneous virtual conferences, each of which may require translation between two or more different languages. For example, a virtual conference between participants from multiple countries may require translating each participant's speech into multiple different languages. Thus, the amount of translation required can be very significant across these various simultaneous conferences. However, the number of translations required at any given time may be difficult to predict. Thus, predicting how to allocate resources for translations can also be difficult.

[0012] To help address this difficulty, a virtual conference provider may maintain a number of servers, each of which has software that may be instantiated

to translate from a particular source language (or multiple different source languages) into one or more target languages. Each instance of the translation software may then be assigned to handle translation functionality for a virtual conference. Thus, a solution may be to receive a request from a participant in a conference, instantiate a new translation software instance, assign it to the conference, and begin providing translation services. However, such a strategy may result in significant lag between a request for translations and translations being made available.

[0013] Another approach may be to maintain a pool of instantiated translation services that may be allocated on demand, but that remain idle until they are allocated. This may significantly reduce the start-up lag for a requested translation. However, because it is difficult to predict how many such translation services may be required at any given time, it remains a problem of maintaining a reasonable number of translation services without wasting energy and processing resources maintaining translation services that are unlikely to be allocated.

[0014] An example system for providing real-time translation during virtual conferences may instead dynamically instantiate or terminate translation services over time based on various factors, such as a number of on-going conferences, whether scheduled conferences include participants from different time zones or different countries, and whether the local time in a particular country (or countries) is within conventional business hours or not. Based on these factors, the system can establish thresholds at which to instantiate an additional pool of translation services or at which to terminate some number of translation services to release processing resources for other activities.

[0015] Such strategies can provide on-demand translation in real-time for participants in a virtual conference while providing low-latency between the request for translation services and when those services are provided, all while providing real-time translation of speech within the virtual conference to the various users.

[0016] This illustrative example is given to introduce the reader to the general subject matter discussed herein and the disclosure is not limited to this example. The following sections describe various additional non-limiting examples and examples of providing real-time translation during virtual conferences.

[0017] Referring now to Figure 1, Figure 1 shows an example system 100 that provides videoconferencing functionality to various client devices. The system 100 includes a virtual conference provider 110 that is connected to multiple communication networks 120, 130, through which various client devices 140-180 can participate in virtual conferences hosted by the virtual conference provider 110. For example, the virtual conference provider 120 can be located within a private network to provide virtual conferencing services to devices within the private network, or it can be connected to a public network, e.g., the internet, so it may be accessed by anyone. Some examples may even provide a hybrid model in which a virtual conference provider 120 may supply components to enable a private organization to host private internal virtual conferences or to connect its system to the virtual conference provider 120 over a public network.

[0018] The system optionally also includes one or more user identity providers, e.g., user identity provider 115, which can provide user identity services to users of the client devices 140-160 and may authenticate user identities of one or more users to the virtual conference provider 110. In this example, the user identity provider 115 is operated by a different entity than the virtual conference provider 110, though in some examples, they may be the same entity.

[0019] Virtual conference provider 110 allows clients to create videoconference meetings (or "meetings") and invite others to participate in those meetings as well as perform other related functionality, such as recording the meetings, generating transcripts from meeting audio, manage user functionality in the meetings, enable text messaging during the meetings, create and manage breakout rooms from the main meeting, etc. Figure 2, described below, provides a more detailed description of the architecture and functionality of the virtual conference provider 110.

[0020] Meetings in this example virtual conference provider 110 are provided in virtual "rooms" to which participants are connected. The room in this context is a construct provided by a server that provides a common point at which the various video and audio data is received before being multiplexed and provided to the various participants. While a "room" is the label for this concept in this disclosure, any suitable functionality that enables multiple participants to participate in a

common videoconference may be used. Further, in some examples, and as alluded to above, a meeting may also have "breakout" rooms. Such breakout rooms may also be rooms that are associated with a "main" videoconference room. Thus, participants in the main videoconference room may exit the room into a breakout room, e.g., to discuss a particular topic, before returning to the main room. The breakout rooms in this example are discrete meetings that are associated with the meeting in the main room. However, to join a breakout room, a participant must first enter the main room. A room may have any number of associated breakout rooms according to various examples.

[0021] To create a meeting with the virtual conference provider 110, a user may contact the virtual conference provider 110 using a client device 140-180 and select an option to create a new meeting. Such an option may be provided in a webpage accessed by a client device 140-160 or client application executed by a client device 140-160. For telephony devices, the user may be presented with an audio menu that they may navigate by pressing numeric buttons on their telephony device. To create the meeting, the virtual conference provider 110 may prompt the user for certain information, such as a date, time, and duration for the meeting, a number of participants, a type of encryption to use, whether the meeting is confidential or open to the public, etc. After receiving the various meeting settings, the virtual conference provider may create a record for the meeting and generate a meeting identifier and, in some examples, a corresponding meeting password or passcode (or other authentication information), all of which meeting information is provided to the meeting host.

[0022] After receiving the meeting information, the user may distribute the meeting information to one or more users to invite them to the meeting. To begin the meeting at the scheduled time (or immediately, if the meeting was set for an immediate start), the host provides the meeting identifier and, if applicable, corresponding authentication information (e.g., a password or passcode). The virtual conference system then initiates the meeting and may admit users to the meeting. Depending on the options set for the meeting, the users may be admitted immediately upon providing the appropriate meeting identifier (and authentication information, as appropriate), even if the host has not yet arrived, or the users may

be presented with information indicating the that meeting has not yet started or the host may be required to specifically admit one or more of the users.

[0023] During the meeting, the participants may employ their client devices 140-180 to capture audio or video information and stream that information to the virtual conference provider 110. They also receive audio or video information from the virtual conference provider 210, which is displayed by the respective client device 140 to enable the various users to participate in the meeting.

[0024] At the end of the meeting, the host may select an option to terminate the meeting, or it may terminate automatically at a scheduled end time or after a predetermined duration. When the meeting terminates, the various participants are disconnected from the meeting and they will no longer receive audio or video streams for the meeting (and will stop transmitting audio or video streams). The virtual conference provider 110 may also invalidate the meeting information, such as the meeting identifier or password/passcode.

[0025] To provide such functionality, one or more client devices 140-180 may communicate with the virtual conference provider 110 using one or more communication networks, such as network 120 or the public switched telephone network ("PSTN") 130. The client devices 140-180 may be any suitable computing or communications device that have audio or video capability. For example, client devices 140-160 may be conventional computing devices, such as desktop or laptop computers having processors and computer-readable media, connected to the virtual conference provider 110 using the internet or other suitable computer network. Suitable networks include the internet, any local area network ("LAN"), metro area network ("MAN"), wide area network ("WAN"), cellular network (e.g., 3G, 4G, 4G LTE, 5G, etc.), or any combination of these. Other types of computing devices may be used instead or as well, such as tablets, smartphones, and dedicated video conferencing equipment. Each of these devices may provide both audio and video capabilities and may enable one or more users to participate in a virtual conference meeting hosted by the virtual conference provider 110.

[0026] In addition to the computing devices discussed above, client devices 140-180 may also include one or more telephony devices, such as cellular telephones (e.g., cellular telephone 170), internet protocol ("IP") phones (e.g., telephone 180), or

conventional telephones. Such telephony devices may allow a user to make conventional telephone calls to other telephony devices using the PSTN, including the virtual conference provider 110. It should be appreciated that certain computing devices may also provide telephony functionality and may operate as telephony devices. For example, smartphones typically provide cellular telephone capabilities and thus may operate as telephony devices in the example system 100 shown in Figure 1. In addition, conventional computing devices may execute software to enable telephony functionality, which may allow the user to make and receive phone calls, e.g., using a headset and microphone. Such software may communicate with a PSTN gateway to route the call from a computer network to the PSTN. Thus, telephony devices encompass any devices that can make conventional telephone calls and is not limited solely to dedicated telephony devices like conventional telephones.

[0027] Referring again to client devices 140-160, these devices 140-160 contact the virtual conference provider 110 using network 120 and may provide information to the virtual conference provider 110 to access functionality provided by the virtual conference provider 110, such as access to create new meetings or join existing meetings. To do so, the client devices 140-160 may provide user identification information, meeting identifiers, meeting passwords or passcodes, etc. In examples that employ a user identity provider 115, a client device, e.g., client devices 140-160, may operate in conjunction with a user identity provider 115 to provide user identification information or other user information to the virtual conference provider 110.

[0028] A user identity provider 115 may be any entity trusted by the virtual conference provider 110 that can help identify a user to the virtual conference provider 110. For example, a trusted entity may be a server operated by a business or other organization and with whom the user has established their identity, such as an employer or trusted third-party. The user may sign into the user identity provider 115, such as by providing a username and password, to access their identity at the user identity provider 115. The identity, in this sense, is information established and maintained at the user identity provider 115 that can be used to identify a particular user, irrespective of the client device they may be using. An example of an identity may be an email account established at the user identity

provider 110 by the user and secured by a password or additional security features, such as biometric authentication, two-factor authentication, etc. However, identities may be distinct from functionality such as email. For example, a health care provider may establish identities for its patients. And while such identities may have associated email accounts, the identity is distinct from those email accounts. Thus, a user's "identity" relates to a secure, verified set of information that is tied to a particular user and should be accessible only by that user. By accessing the identity, the associated user may then verify themselves to other computing devices or services, such as the virtual conference provider 110.

[0029] When the user accesses the virtual conference provider 110 using a client device, the virtual conference provider 110 communicates with the user identity provider 115 using information provided by the user to verify the user's identity. For example, the user may provide a username or cryptographic signature associated with a user identity provider 115. The user identity provider 115 then either confirms the user's identity or denies the request. Based on this response, the virtual conference provider 110 either provides or denies access to its services, respectively.

[0030] For telephony devices, e.g., client devices 170-180, the user may place a telephone call to the virtual conference provider 110 to access virtual conference services. After the call is answered, the user may provide information regarding a virtual conference meeting, e.g., a meeting identifier ("ID"), a passcode or password, etc., to allow the telephony device to join the meeting and participate using audio devices of the telephony device, e.g., microphone(s) and speaker(s), even if video capabilities are not provided by the telephony device.

[0031] Because telephony devices typically have more limited functionality than conventional computing devices, they may be unable to provide certain information to the virtual conference provider 110. For example, telephony devices may be unable to provide user identification information to identify the telephony device or the user to the virtual conference provider 110. Thus, the virtual conference provider 110 may provide more limited functionality to such telephony devices. For example, the user may be permitted to join a meeting after providing meeting information, e.g., a meeting identifier and passcode, but they may be

identified only as an anonymous participant in the meeting. This may restrict their ability to interact with the meetings in some examples, such as by limiting their ability to speak in the meeting, hear or view certain content shared during the meeting, or access other meeting functionality, such as joining breakout rooms or engaging in text chat with other participants in the meeting.

[0032] It should be appreciated that users may choose to participate in meetings anonymously and decline to provide user identification information to the virtual conference provider 110, even in cases where the user has an authenticated identity and employs a client device capable of identifying the user to the virtual conference provider 110. The virtual conference provider 110 may determine whether to allow such anonymous users to use services provided by the virtual conference provider 110. Anonymous users, regardless of the reason for anonymity, may be restricted as discussed above with respect to users employing telephony devices, and in some cases may be prevented from accessing certain meetings or other services, or may be entirely prevented from accessing the virtual conference provider 110.

[0033] Referring again to virtual conference provider 110, in some examples, it may allow client devices 140-160 to encrypt their respective video and audio streams to help improve privacy in their meetings. Encryption may be provided between the client devices 140-160 and the virtual conference provider 110 or it may be provided in an end-to-end configuration where multimedia streams transmitted by the client devices 140-160 are not decrypted until they are received by another client device 140-160 participating in the meeting. Encryption may also be provided during only a portion of a communication, for example encryption may be used for otherwise unencrypted communications that cross international borders.

[0034] Client-to-server encryption may be used to secure the communications between the client devices 140-160 and the virtual conference provider 110, while allowing the virtual conference provider 110 to access the decrypted multimedia streams to perform certain processing, such as recording the meeting for the participants or generating transcripts of the meeting for the participants. End-to-end encryption may be used to keep the meeting entirely private to the participants without any worry about a virtual conference provider 110 having access to the

substance of the meeting. Any suitable encryption methodology may be employed, including key-pair encryption of the streams. For example, to provide end-to-end encryption, the meeting host's client device may obtain public keys for each of the other client devices participating in the meeting and securely exchange a set of keys to encrypt and decrypt multimedia content transmitted during the meeting. Thus the client devices 140-160 may securely communicate with each other during the meeting. Further, in some examples, certain types of encryption may be limited by the types of devices participating in the meeting. For example, telephony devices may lack the ability to encrypt and decrypt multimedia streams. Thus, while encrypting the multimedia streams may be desirable in many instances, it is not required as it may prevent some users from participating in a meeting.

[0035] By using the example system shown in Figure 1, users can create and participate in meetings using their respective client devices 140-180 via the virtual conference provider 110. Further, such a system enables users to use a wide variety of different client devices 140-180 from traditional standards-based video conferencing hardware to dedicated video conferencing equipment to laptop or desktop computers to handheld devices to legacy telephony devices, etc.

[0036] Referring now to Figure 2, Figure 2 shows an example system 200 in which a virtual conference provider 210 provides videoconferencing functionality to various client devices 220-250. The client devices 220-250 include two conventional computing devices 220-230, dedicated equipment for a virtual conference room 240, and a telephony device 250. Each client device 220-250 communicates with the virtual conference provider 210 over a communications network, such as the internet for client devices 220-240 or the PSTN for client device 250, generally as described above with respect to Figure 1. The virtual conference provider 210 is also in communication with one or more user identity providers 215, which can authenticate various users to the virtual conference provider 210 generally as described above with respect to Figure 1.

[0037] In this example, the virtual conference provider 210 employs multiple different servers (or groups of servers) to provide different aspects of virtual conference functionality, thereby enabling the various client devices to create and participate in virtual conference meetings. The virtual conference provider 210 uses

one or more real-time media servers 212, one or more network services servers 214, one or more video room gateways 216, and one or more telephony gateways 218. Each of these servers 212-218 is connected to one or more communications networks to enable them to collectively provide access to and participation in one or more virtual conference meetings to the client devices 220-250.

[0038] The real-time media servers 212 provide multiplexed multimedia streams to meeting participants, such as the client devices 220-250 shown in Figure 2. While video and audio streams typically originate at the respective client devices, they are transmitted from the client devices 220-250 to the virtual conference provider 210 via one or more networks where they are received by the real-time media servers 212. The real-time media servers 212 determine which protocol is optimal based on, for example, proxy settings and the presence of firewalls, etc. For example, the client device might select among UDP, TCP, TLS, or HTTPS for audio and video and UDP for content screen sharing.

[0039] The real-time media servers 212 then multiplex the various video and audio streams based on the target client device and communicate multiplexed streams to each client device. For example, the real-time media servers 212 receive audio and video streams from client devices 220-240 and only an audio stream from client device 250. The real-time media servers 212 then multiplex the streams received from devices 230-250 and provide the multiplexed streams to client device 220. The real-time media servers 212 are adaptive, for example, reacting to real-time network and client changes, in how they provide these streams. For example, the real-time media servers 212 may monitor parameters such as a client's bandwidth CPU usage, memory and network I/O as well as network parameters such as packet loss, latency and jitter to determine how to modify the way in which streams are provided.

[0040] The client device 220 receives the stream, performs any decryption, decoding, and demultiplexing on the received streams, and then outputs the audio and video using the client device's video and audio devices. In this example, the real-time media servers do not multiplex client device 220's own video and audio feeds when transmitting streams to it. Instead each client device 220-250 only receives multimedia streams from other client devices 220-250. For telephony devices that

lack video capabilities, e.g., client device 250, the real-time media servers 212 only deliver multiplex audio streams. The client device 220 may receive multiple streams for a particular communication, allowing the client device 220 to switch between streams to provide a higher quality of service.

[0041] In addition to multiplexing multimedia streams, the real-time media servers 212 may also decrypt incoming multimedia stream in some examples. As discussed above, multimedia streams may be encrypted between the client devices 220-250 and the virtual conference system 210. In some such examples, the real-time media servers 212 may decrypt incoming multimedia streams, multiplex the multimedia streams appropriately for the various clients, and encrypt the multiplexed streams for transmission.

[0042] In some examples, to provide multiplexed streams, the virtual conference provider 210 may receive multimedia streams from the various participants and publish those streams to the various participants to subscribe to and receive. Thus, the virtual conference provider 210 notifies a client device, e.g., client device 220, about various multimedia streams available from the other client devices 230-250, and the client device 220 can select which multimedia stream(s) to subscribe to and receive. In some examples, the virtual conference provider 210 may provide to each client device the available streams from the other client devices, but from the respective client device itself, though in other examples it may provide all available streams to all available client devices. Using such a multiplexing technique, the virtual conference provider 210 may enable multiple different streams of varying quality, thereby allowing client devices to change streams in real-time as needed, e.g., based on network bandwidth, latency, etc.

[0043] As mentioned above with respect to Figure 1, the virtual conference provider 210 may provide certain functionality with respect to unencrypted multimedia streams at a user's request. For example, the meeting host may be able to request that the meeting be recorded or that a transcript of the audio streams be prepared, which may then be performed by the real-time media servers 212 using the decrypted multimedia streams, or the recording or transcription functionality may be off-loaded to a dedicated server (or servers), e.g., cloud recording servers, for recording the audio and video streams. In some examples, the virtual conference

provider 210 may allow a meeting participant to notify it of inappropriate behavior or content in a meeting. Such a notification may trigger the real-time media servers to 212 record a portion of the meeting for review by the virtual conference provider 210. Still other functionality may be implemented to take actions based on the decrypted multimedia streams at the virtual conference provider, such as monitoring video or audio quality, adjusting or changing media encoding mechanisms, etc.

[0044] It should be appreciated that multiple real-time media servers 212 may be involved in communicating data for a single meeting and multimedia streams may be routed through multiple different real-time media servers 212. In addition, the various real-time media servers 212 may not be co-located, but instead may be located at multiple different geographic locations, which may enable high-quality communications between clients that are dispersed over wide geographic areas, such as being located in different countries or on different continents. Further, in some examples, one or more of these servers may be co-located on a client's premises, e.g., at a business or other organization. For example, different geographic regions may each have one or more real-time media servers 212 to enable client devices in the same geographic region to have a high-quality connection into the virtual conference provider 210 via local servers 212 to send and receive multimedia streams, rather than connecting to a real-time media server located in a different country or on a different continent. The local real-time media servers 212 may then communicate with physically distant servers using high-speed network infrastructure, e.g., internet backbone network(s), that otherwise might not be directly available to client devices 220-250 themselves. Thus, routing multimedia streams may be distributed throughout the virtual conference system 210 and across many different real-time media servers 212.

[0045] Turning to the network services servers 214, these servers 214 provide administrative functionality to enable client devices to create or participate in meetings, send meeting invitations, create or manage user accounts or subscriptions, and other related functionality. Further, these servers may be configured to perform different functionalities or to operate at different levels of a hierarchy, e.g., for specific regions or localities, to manage portions of the virtual conference provider

under a supervisory set of servers. When a client device 220-250 accesses the virtual conference provider 210, it will typically communicate with one or more network services servers 214 to access their account or to participate in a meeting.

[0046] When a client device 220-250 first contacts the virtual conference provider 210 in this example, it is routed to a network services server 214. The client device may then provide access credentials for a user, e.g., a username and password or single sign-on credentials, to gain authenticated access to the virtual conference provider 210. This process may involve the network services servers 214 contacting a user identity provider 215 to verify the provided credentials. Once the user's credentials have been accepted, the client device 220-250 may perform administrative functionality, like updating user account information, if the user has an identity with the virtual conference provider 210, or scheduling a new meeting, by interacting with the network services servers 214.

[0047] In some examples, users may access the virtual conference provider 210 anonymously. When communicating anonymously, a client device 220-250 may communicate with one or more network services servers 214 but only provide information to create or join a meeting, depending on what features the virtual conference provider allows for anonymous users. For example, an anonymous user may access the virtual conference provider using client 220 and provide a meeting ID and passcode. The network services server 214 may use the meeting ID to identify an upcoming or on-going meeting and verify the passcode is correct for the meeting ID. After doing so, the network services server(s) 214 may then communicate information to the client device 220 to enable the client device 220 to join the meeting and communicate with appropriate real-time media servers 212.

[0048] In cases where a user wishes to schedule a meeting, the user (anonymous or authenticated) may select an option to schedule a new meeting and may then select various meeting options, such as the date and time for the meeting, the duration for the meeting, a type of encryption to be used, one or more users to invite, privacy controls (e.g., not allowing anonymous users, preventing screen sharing, manually authorize admission to the meeting, etc.), meeting recording options, etc. The network services servers 214 may then create and store a meeting record for the scheduled meeting. When the scheduled meeting time arrives (or

within a threshold period of time in advance), the network services server(s) 214 may accept requests to join the meeting from various users.

[0049] To handle requests to join a meeting, the network services server(s) 214 may receive meeting information, such as a meeting ID and passcode, from one or more client devices 220-250. The network services server(s) 214 locate a meeting record corresponding to the provided meeting ID and then confirm whether the scheduled start time for the meeting has arrived, whether the meeting host has started the meeting, and whether the passcode matches the passcode in the meeting record. If the request is made by the host, the network services server(s) 214 activates the meeting and connects the host to a real-time media server 212 to enable the host to begin sending and receiving multimedia streams.

[0050] Once the host has started the meeting, subsequent users requesting access will be admitted to the meeting if the meeting record is located and the passcode matches the passcode supplied by the requesting client device 220-250. In some examples additional access controls may be used as well. But if the network services server(s) 214 determines to admit the requesting client device 220-250 to the meeting, the network services server 214 identifies a real-time media server 212 to handle multimedia streams to and from the requesting client device 220-250 and provides information to the client device 220-250 to connect to the identified real-time media server 212. Additional client devices 220-250 may be added to the meeting as they request access through the network services server(s) 214.

[0051] After joining a meeting, client devices will send and receive multimedia streams via the real-time media servers 212, but they may also communicate with the network services servers 214 as needed during meetings. For example, if the meeting host leaves the meeting, the network services server(s) 214 may appoint another user as the new meeting host and assign host administrative privileges to that user. Hosts may have administrative privileges to allow them to manage their meetings, such as by enabling or disabling screen sharing, muting or removing users from the meeting, creating sub-meetings or “break-out” rooms, recording meetings, etc. Such functionality may be managed by the network services server(s) 214.

[0052] For example, if a host wishes to remove a user from a meeting, they may identify the user and issue a command through a user interface on their client device. The command may be sent to a network services server 214, which may then disconnect the identified user from the corresponding real-time media server 212. If the host wishes to create a break-out room for one or more meeting participants to join, such a command may also be handled by a network services server 214, which may create a new meeting record corresponding to the break-out room and then connect one or more meeting participants to the break-out room similarly to how it originally admitted the participants to the meeting itself.

[0053] In addition to creating and administering on-going meetings, the network services server(s) 214 may also be responsible for closing and tearing-down meetings once they have completed. For example, the meeting host may issue a command to end an on-going meeting, which is sent to a network services server 214. The network services server 214 may then remove any remaining participants from the meeting, communicate with one or more real time media servers 212 to stop streaming audio and video for the meeting, and deactivate, e.g., by deleting a corresponding passcode for the meeting from the meeting record, or delete the meeting record(s) corresponding to the meeting. Thus, if a user later attempts to access the meeting, the network services server(s) 214 may deny the request.

[0054] Depending on the functionality provided by the virtual conference provider, the network services server(s) 214 may provide additional functionality, such as by providing private meeting capabilities for organizations, special types of meetings (e.g., webinars), etc. Such functionality may be provided according to various examples of virtual conference providers according to this description.

[0055] Referring now to the video room gateway servers 216, these servers 216 provide an interface between dedicated video conferencing hardware, such as may be used in dedicated video conferencing rooms. Such video conferencing hardware may include one or more cameras and microphones and a computing device designed to receive video and audio streams from each of the cameras and microphones and connect with the virtual conference provider 210. For example, the video conferencing hardware may be provided by the virtual conference provider to

one or more of its subscribers, which may provide access credentials to the video conferencing hardware to use to connect to the virtual conference provider 210.

[0056] The video room gateway servers 216 provide specialized authentication and communication with the dedicated video conferencing hardware that may not be available to other client devices 220-230, 250. For example, the video conferencing hardware may register with the virtual conference provider 210 when it is first installed and the video room gateway servers 216 may authenticate the video conferencing hardware using such registration as well as information provided to the video room gateway server(s) 216 when dedicated video conferencing hardware connects to it, such as device ID information, subscriber information, hardware capabilities, hardware version information etc. Upon receiving such information and authenticating the dedicated video conferencing hardware, the video room gateway server(s) 216 may interact with the network services servers 214 and real-time media servers 212 to allow the video conferencing hardware to create or join meetings hosted by the virtual conference provider 210.

[0057] Referring now to the telephony gateway servers 218, these servers 218 enable and facilitate telephony devices' participation in meetings hosed by the virtual conference provider 210. Because telephony devices communicate using the PSTN and not using computer networking protocols, such as TCP/IP, the telephony gateway servers 218 act as an interface that converts between the PSTN and the networking system used by the virtual conference provider 210.

[0058] For example, if a user uses a telephony device to connect to a meeting, they may dial a phone number corresponding to one of the virtual conference provider's telephony gateway servers 218. The telephony gateway server 218 will answer the call and generate audio messages requesting information from the user, such as a meeting ID and passcode. The user may enter such information using buttons on the telephony device, e.g., by sending dual-tone multi-frequency ("DTMF") audio signals to the telephony gateway server 218. The telephony gateway server 218 determines the numbers or letters entered by the user and provides the meeting ID and passcode information to the network services servers 214, along with a request to join or start the meeting, generally as described above. Once the

telephony client device 250 has been accepted into a meeting, the telephony gateway server 218 is instead joined to the meeting on the telephony device's behalf.

[0059] After joining the meeting, the telephony gateway server 218 receives an audio stream from the telephony device and provides it to the corresponding real-time media server 212, and receives audio streams from the real-time media server 212, decodes them, and provides the decoded audio to the telephony device. Thus, the telephony gateway servers 218 operate essentially as client devices, while the telephony device operates largely as an input/output device, e.g., a microphone and speaker, for the corresponding telephony gateway server 218, thereby enabling the user of the telephony device to participate in the meeting despite not using a computing device or video.

[0060] It should be appreciated that the components of the virtual conference provider 210 discussed above are merely examples of such devices and an example architecture. Some virtual conference providers may provide more or less functionality than described above and may not separate functionality into different types of servers as discussed above. Instead, any suitable servers and network architectures may be used according to different examples.

[0061] Referring now to Figure 3, Figure 3 shows an example system 300 for providing real-time translation during virtual conferences. The system 300 includes a virtual conference provider 310, which can be connected to multiple client device 330, 340a-n via one or more intervening communication networks 320. In this example, the communications network 320 is the internet, however, any suitable communications network or combination of communications network may be employed, including LANs (e.g., within a corporate private LAN), WANs, etc.

[0062] Each client device 330, 340a-*n* executes virtual conference software, which connects to the virtual conference provider 310 and joins a meeting. During the meeting, the various participants (using virtual conference software or "client software" at their respective client devices 330, 340a-*n*) are able to interact with each other to conduct the meeting, such as by viewing video feeds and hearing audio feeds from other participants, and by capturing and transmitting video and audio of themselves.

[0063] The virtual conference provider 310 operates a number of servers 312 that can provide transcription and translation functionality. Transcription functionality is provided by one or more transcription processes 314 that can be executed and allocated to virtual conferences hosted by the virtual conference provider 310. The transcription processes 314 in this example employ a respective deep neural networks ("DNN") to convert speech to text in one or more languages; however, any suitable machine learning model or speech-to-text conversion functionality may be employed. Similarly, translation functionality is provided by one or more translation processes 316 that can be executed and allocated to virtual conferences hosted by the virtual conference provider 310. As with the transcription processes 314, the translation processors 316 in this example also employ a respective DNN to translate text from one or more source languages to one or more target languages; however, any suitable machine learning model or text translation functionality may be employed.

[0064] Client device 330, 340a-n may join virtual conferences hosted by the virtual conference provider 310 by connecting to the virtual conferences provider and joining a desired virtual conference, generally as discussed above with respect to Figures 1-2. Once the participants have joined the conference, they may interact with each other by exchanging audio and video feeds. However, if different participants do not speak a common language with each other, they can request translation services from the virtual conference provider 310.

[0065] To request translation services, a participant may select an option within their client software to enable translations. They may then select one or more source languages, one or more participants, and a target language for the translation(s). The client software then sends a request to the virtual conference provider 310 for the selected translation services.

[0066] After receiving a request for translation services, the virtual conference provider 310 allocates one or more transcription processes 314 to the virtual conference, with one transcription process 314 allocated for each selected source language. In this example, each transcription process 314 is configured to generate a transcript in multiple input languages, however, in some cases, a transcription process 314 may only handle a single input language. Further, in some

examples, one or more transcription processes 314 may be configured to handle more than one input audio stream, and thus may be able to convert speech into text simultaneously for each of the input streams. Thus, appropriate transcription processes may be selected depending on the capabilities of the available transcription processes.

[0067] Similarly, the virtual conference provider 310 allocates one or more translation processes 316, depending on the selections made by the participant. As with the transcription processes, each translation process 316 is configured to translate from multiple sources languages into a specific target language. Though other examples may translate from any one of multiple source languages into any one of multiple target languages, or they may be specific to a single source language and a single target language. However, based on its configuration, the virtual conference provider 310 allocates the appropriate translation process(es) 316 to the virtual conference.

[0068] The transcription process(es) 314 receive audio streams from the virtual conference and each generate a textual representation of spoken words in the source language. The textual representation may then be stored as a transcript of the call, but it is also provided to the translation process(es) 316 to translate into a textual representation in the target language(s). The translated text is then provided to the requesting client device(s) 330, 340a-n and provided as closed captions or in a separate display area of the client software's GUI to enable the participant to view the translation.

[0069] To allow the translation to occur in real-time, the transcription process 314 streams its output to the translation process 316, which translates as it receives the text. However, simply streaming transcribed text may not provide an accurate or low-latency translation due to differences in sentence structure between different languages, e.g., some languages typically place verbs near the end of the sentence while others typically place them near the beginning, or due to a word or phrase appearing to have one meaning standing alone, but a different meaning once more fulsome context is available, e.g., from subsequently spoken words.

[0070] To accommodate such issues, the translation process 316 may output translated text as soon as it is available on a word-by-word basis, however, it may

retain memory of recently translated text and, as new text arrives from a transcription process, it may update its translation based on additional context or other parts of speech that were previously missing. The updated translation may be provided to the client device and may be flagged as an updated translation, which then allows the client device to replace previously displayed translated text with the updated translation. Thus, for each client device that requests translation from a source language (or multiple source languages) to a target language, the virtual conference provider 310 will allocate the appropriate transcription and translation processes 314, 316.

[0071] Once the virtual conference has concluded, the virtual conference provider 310 de-allocates the allocated transcription and translation processes 314, 316 from the virtual conference and returns them to the pool of available, but idle transcription and translation processes 314, 316, making them available to be allocated to other virtual conferences or for termination if the virtual conference provider 310 determines it has too many idle transcription or translation processes 314, 316.

[0072] Referring now to Figure 4, Figure 4 shows an example data flow diagram for a system that provides real-time translation during virtual conferences. The system includes a virtual conference provider 410 that is hosting a virtual conference between multiple client devices, such as client device 330, 340a-n. The virtual conference provider 410 has received a request to translate one or more audio streams in a source language to a target language. During the virtual conference, the virtual conference provider 410 receives audio streams 450 from the various client devices and identifies one or more audio streams in the first language, such as by receiving an identification of a participant within the translation request or by performing language identification on the received audio streams.

[0073] The audio stream(s) 450 to be translated are then provided to a transcription process 414, which generates one or more words 452 that contain a textual representation of speech contained in the audio stream(s). The words 452 are then provided to one or more allocated translation processes 416. If multiple participants have requested translation of the same source language into different target languages, multiple different translation processes 416 may be allocated,

though a single transcription process 414 may generate a transcript and provide it to the respective transcription processes 416. Similarly, any number of source and target languages may be employed, which would involve using a suitable set of transcription and translation processes 414, 416 based on the selected languages. After generating the words 452, they may also be stored as a transcript 454 in a storage medium and later provided to one or more participants, if a transcript of the virtual conference has been requested. Otherwise, the words 452 may be discarded by the translation process 416 once they are no longer needed for translation.

[0074] As discussed above, the words 452 and translation 456 may be generated in real-time, and thus, the words 452 may not include the entirety of the speech from the length of the conference, but may only include individual words or groups of words as they are received within the audio stream(s). Similarly, the translation 456 may be output one or more words at a time. Further, as discussed above, if the translation process 416 updates a translation based on subsequently received words 452, it may output a translation 456 including an indication that the translation 456 updates previously output translated words. The translation 456 may also include additional information to specify which portion of the previously output translated words are to be replaced with the new translation 456.

[0075] Further, and in addition to translation audio streams, the virtual conference provider 410 may also translate chat messages sent using chat functionality within the software client. Since chat messages are already in a textual form, they may be provided directly to the appropriate translation process 416 and the translation 456 corresponding to the chat messages 460 may be generated and output to any requesting participants.

[0076] While the example shown in Figure 4 illustrates only a single transcription process 414 and a single translation process 416, any number of such processes 414, 416 may operate simultaneously for a single virtual conference. Further the virtual conference provider 410 may host multiple concurrent virtual conferences, each of which may be allocated one or more transcription and translation processes 414, 416.

[0077] Referring now to Figures 5A-5B, Figure 5A illustrates an example GUI 500 for a software client that can interact with a system for providing real-time

translation during virtual conferences. A client device, e.g., client device 330 or client devices 340a-n, executes a software client as discussed above, which in turn displays the GUI 500 on the client device's display. In this example, the GUI 500 includes a speaker view window 502 that presents the current speaker in the virtual conference. Above the speaker view window 502 are smaller participant windows 504, which allow the participant to view some of the other participants in the virtual conference, as well as controls ("<" and ">") to let the host scroll to view other participants in the virtual conference. On the right side of the GUI 500 is a chat window 540 within which the participants may exchange chat messages.

[0078] Beneath the speaker view window 502 are a number of interactive elements 510-530 to allow the participant to interact with the virtual conference software. Controls 510-512 may allow the participant to toggle on or off audio or video streams captured by a microphone or camera connected to the client device. Control 520 allows the participant to view any other participants in the virtual conference with the participant, while control 522 allows the participant to send text messages to other participants, whether to specific participants or to the entire meeting. Control 524 allows the participant to share content from their client device. Control 526 allows the participant toggle recording of the meeting, and control 528 allows the user to select an option to join a breakout room. Control 530 allows a user to launch an app within the virtual conferencing software, such as to access content to share with other participants in the virtual conference.

[0079] In addition to controls 510-530, a translation button 532 is provided by the GUI 500 which allows the user to request translations of audio within a virtual conference. When the translation button 532 is selected, the user may be presented with the option to select a source or target language or one or more participants or the chat window 540 to translate. In some examples, the virtual conference provider may automatically detect languages, or participants may establish the language(s) they understand or speak within a user profile or when accessing the virtual conference. Thus, in some examples the user may not need to select a source or target language. For example, if each user's language is automatically detected or otherwise provided, such as by a user profile, the virtual conference provider may select the user's language as the target language and any

participants who are either not speaking the user's language or do not have the user's language set in their profile, the virtual conference provider may automatically select source and target languages for translation functionality.

[0080] During the normal course of a virtual conference, the user interacts with the virtual conferencing application and other participants via the GUI 500. And if translation functionality is desired, the user may select the translation button 532.

[0081] Figure 5B illustrates the GUI 500 after the user has pressed the translation button 532, which has darkened to indicate the translation functionality is active. In addition, translation notifications 534 have been overlaid on participant windows within the conference whose audio or chat messages are being translated. In some examples, a user may interact with a translation notification 534 to determine a source language for the participant, such as in an example where the virtual conference provider is able automatically determine a source language for each speaker.

[0082] In addition to providing the translation notifications 534, the GUI 500 provides a translation overlay 550 on the speaker window 502, which provides the output from the translation process(es) that have been allocated to the virtual conference and that correspond to the translations requested by the user. The translation overlay 550 may provide scrolling text corresponding to the real-time translation of the audio streams and further may identify the speaker corresponding to each translated audio stream. Thus, the user may read along with the translation to understand what is being said. In some examples the translation overlay 550 may also identify the speaker or identity that uttered the translated text. Similarly, if participants whose audio streams are being translated make use of the chat functionality, the chat window 540 may present the translated text corresponding to the original chat messages in the source language. Some examples may show both the untranslated and the translated chat messages for the user. Further, some examples may provide a full transcript view for the entire transcript of the virtual conference, in the source or target languages (or both), such as in an additional text window adjacent to the chat window 540 and may also include the identities of the various speakers. The user may then scroll through the transcript as desired during

the conference. Thus, by selecting the translate button, the user is able to quickly and easily enable translation functionality within a virtual conference and see the translations in real-time during the virtual conference.

[0083] Referring now to Figure 6, Figure 6 shows an example method 600 for providing real-time translation during virtual conferences. The example method 600 will be described with respect to the system shown in Figure 4; however any suitable system according to this disclosure may be employed, including the example systems 100-300 shown in Figures 1-3.

[0084] At block 610, the virtual conference provider 310 executes a plurality of transcription and translation processes. In this example, each executing transcription and translation process is associated with one or more of a number of supported source languages, while each executing translation process is able to translate from one or more source languages to one or more target languages. In this example, transcription and translation processes corresponding to multiple different source languages are executed, with the translation processes corresponding to multiple different target languages. By providing transcription and translation processes associated with multiple different source and target languages, the virtual conference provider 310 is able to handle translation requests for various different source and target languages.

[0085] At block 620, the virtual conference provider 310 hosts multiple virtual conferences over a period of time, which may include concurrently hosting multiple virtual conferences, generally as described above with respect to Figures 1 and 2.

[0086] At block 630, the virtual conference provider 310 receives requests from one or more client devices to translate audio streams within their respective virtual conferences. For example, as discussed above with respect to Figure 5A, a user may press a translation button 532 to request translation. The request may specify a source language, a target language, or one or more participants. For example, when the user presses the translation button 532, they may be prompted for source and target languages. They may further be requested to identify one or more participants that are speaking the identified source language. Such

information may be included within the request to enable the virtual conference provider to allocate appropriate transcription and translation processes 414, 416.

[0087] In some examples, however, the request may only specify a request for translation and the virtual conference provider 310 may identify a language spoken by the requesting user and languages spoken by other participants that do not also speak the user's language. The virtual conference provider 310 may thus infer one or more language mismatches between the participants and determine source and target languages based on the mismatch(es). In some examples, the user may only specify a source language, but without specifying a participant. The virtual conference provider 310 may then detect one or more participants speaking the source language and provide their audio streams for translation.

[0088] At block 640, the virtual conference provider 310 allocates one or more transcription processes and one or more translation processes to the respective virtual conferences. As discussed above with respect to Figures 3 and 4, the transcription processes may be allocated from a pool of executing, but idle, transcription processes that are able to perform speech-to-text transcription from an identified source language. Thus, for each identified source language for translation in a particular virtual conference, the virtual conference provider 310 identifies an executing, but idle, transcription process 414 that can process the source language. If available transcription processes 414 are able to process multiple source languages, the virtual conference provider 310 may allocate transcription processes 414 that can process multiple of the identified source languages.

[0089] Similarly, the virtual conference provider 310 allocates translation processes based on the identified source and target languages. If available translation processes 416 are only able to translate from one source language to one target language, suitable translation processes 416 are identified and allocated based on the requested source and target languages for translation. But if available translation processes 416 are able to process translate either multiple source languages into one or more target languages, or one source language into multiple target languages, the virtual conference provider 310 may allocate translation processes 416 that can translate between multiple of the identified source or target languages.

[0090] At block 650, the virtual conference provider 310 provides one or more audio streams from a respective conference to a transcription process 414 allocated to the respective conference according to the source language in the audio stream and the source language(s) that the transcription process 414 is configured to process. As discussed above, the translation request may identify particular participants as speaking particular source languages, or the virtual conference provider 310 may determine the languages spoken by the various participants. It may then identify audio streams from the respective participants and provide them to the transcription process corresponding to the source languages.

[0091] At block 660, the virtual conference provider 310 provides an output from the respective transcription processes 414 for a respective virtual conference to an allocated translation process 416 based on the source and target languages for translation. As discussed above, because the system 300 provides real-time translations, the transcription process 414 in this example outputs individual words or phrases to the translation process 416 as they are generated by the transcription process 414. In addition, if one or more participants in a respective virtual conference have requested a transcript, the transcription process 414 may also generate a transcript, e.g., transcript 454, in the source language(s) within the conference. If multiple source languages are being used, the outputs of the different transcription processes may be merged together into a common transcript, such as by generating and storing timestamps with respect to transcribed speech. The different transcription outputs from different transcription processes 414 may then be merged according to the sequence of timestamps from the various transcriptions.

[0092] At block 670, the virtual conference provider 310 provides an output from the respective translation processes 416 for a respective virtual conference to the requesting client device(s). As discussed above, the output of the translation process 416 may be streamed in real-time to the client device that sent the translation request. Streaming in real-time may involve sending individual words or phrases as they are output by the translation process 416. The translation may be streamed in only the target language, or it may also be streamed in both the source and target language. The manner of providing the streamed output may be

established by the request, by a subsequent request for one or both output streams, or by a configuration of the translation process 416.

[0093] In some examples, the translation process 416 may also output an indicator, such as a flag, if output represents an update to prior output. For example, during translation a particular word may translate to word in the target language, but after receiving additional words from the transcription process 414, the translation process 416 may revise a previously translated word or phrase. It may then output the new translation, an indicator to inform the client device of the change, and an identification of the word or words to be replaced by the new translation. This may enable the system to provide real-time translation, while enabling adjustments to the translation as more context is received.

[0094] In some examples, the system 300 may also store the translation as a translated transcript and provide the translated transcript to the user that requested the translation. Further, in some examples, the translated transcript may be merged with a transcript generated by the transcript process(es), such as based on timestamps as discussed above. However, because the translation may have identical timestamps to the source language transcript, the merged transcript may present the source and translated transcriptions sequentially within the merged transcript. For example, a sentence translated from the source language may be stored in a merged transcript immediately following the transcription of the corresponding source language sentence. Thus, a later viewer of the transcript can see the two (or more) languages adjacent to each other.

[0095] It should be appreciated that the method above may be performed multiple times in sequence or concurrently over a period of time. Further, as new virtual conferences begin and existing virtual conferences end, the method may restart or re-enter any number of times for the new conferences and terminate any number of times as virtual conferences end. Further, while the method was described with respect to multiple concurrent conferences, it should be appreciated that the method could be performed with respect to a single conference and thus, may be executed many times concurrently for individual conferences.

[0096] Referring now to Figure 7, Figure 7 illustrates an example method 700 for providing real-time translation during virtual conferences. The example method

700 will be described with respect to the system shown in Figure 4; however any suitable system according to this disclosure may be employed, including the example systems 100-300 shown in Figures 1-3.

[0097] At block 710, the video conference provider 310 determines a number of available translation or transcription processes 414, 416. In this example, the video conference provider 310 determines a number of translation or transcription processes 414, 416 that are executing, but are idle. As discussed above, an idle translation or transcription process 414, 416 is a process that has been executed by the video conference provider 310, e.g., by a server, but has not been allocated to a virtual conference. As discussed above with respect to Figure 6, as participants in virtual conferences request translation functionality, idle transcription or translation processes 414, 416 are allocated to the corresponding virtual conference, which reduces the number available. Similarly, as virtual conferences end, any transcription or translation processes allocated may be returned to the pool of available, but idle translation and transcription processes.

[0098] It should be appreciated that, because both transcription and translation processes 414, 416 may be configured to process a particular source language or generate a particular target language, the virtual conference provider 310 may determine many different numbers of available transcription and translation processes. For example, the virtual conference provider 310 may maintain individual counts of transcription processes by source language, e.g., individual counts for transcription processes configured to transcribe English, French, Spanish, etc., and individual counts of translation processes by source and target language, e.g., individual counts for translation processes configured to translate from English to French, English to Spanish, etc. For processes that are configured to process multiple source or target languages, counts for each source language or source-to-target language may be incremented as new processes become available or decremented as processes are allocated. Thus, at any given time, the virtual conference provider 310 may maintain many different individual counts.

[0099] At block 720, the video conference provider 310 determines a first threshold associated with one or more of available transcription or translation processes 314, 316. To manage the processing and memory within the various

servers operated by the virtual conference provider 310, it may maintain one or more thresholds corresponding to the various configurations of transcription and translation processes 314, 316 that are available but idle. Each of these thresholds may be established individually and may change over time.

[00100] For example, during the business day in Europe, thresholds for available but idle transcription and translation processes for various European languages may be increased, while thresholds for available but idle transcription processes for various Asian languages may be decreased. Similarly, over the course of a day, the thresholds may be reduced as different countries begin or end their respective business days. The reduction may be in stages, e.g., thresholds may reduce in steps from a highest available threshold to a lowest available threshold, according to different time zones entering or exiting the business day. Further, different thresholds may be used on non-business days, e.g., weekends or holidays. Other factors may further affect thresholds. For example, a number of scheduled webinars and locations of the various of invitees may be used to determine an increase from a standard business-day threshold. For example, large, international business conferences may be scheduled over the course of one or more days, but ordinarily would not factor into a need for transcription or translation functionality. Thus, the virtual conference provider 310 can increase its thresholds when it determines that such events have been scheduled.

[00101] By employing dynamic thresholds the virtual conference provider 310 may be able to maintain a sufficient number of idle transcription and translation processes to ensure low-latency responses to translation requests, while also not inefficiently consuming processing and memory resources when the likelihood of significant numbers of translation requests for a particular source-to-target language combination is low.

[00102] At block 730, the virtual conference provider 310 determines whether one or more of the thresholds has been satisfied. For each threshold that is satisfied, the method 700 proceeds to block 732. For each threshold that is not satisfied, the method 700 proceeds to block 740.

[00103] At block 732, the virtual conference provider 310 instantiates a number of new transcription or translation processes 414, 416 needed to bring the

total number of idle transcription or translation processes 414, 416 of the particular configuration up to the threshold. Thus, the system can determine if too few of a particular type of translation or transcription process 414, 416 are idle and can increase the size of the pool. The method 700 then returns to block 710.

[00104] At block 740, the virtual conference provider 310 determines a second threshold. As discussed above with respect to block 720, may different thresholds may be maintained for various configurations of translation and transcription processes 414, 416.

[00105] At block 750, the virtual conference provider 310 determines whether the second threshold is met for one or more configurations of transcription or translation processes 414, 416. In this case the threshold relate to whether too many transcription or translation processes 414, 416 are idle. If the threshold is satisfied, the method 700 proceeds to block 752. Otherwise, it returns to block 710.

[00106] At block 752, the virtual conference provider 310 terminates one or more transcription or translation processes 414, 416 to reduce the number of idle processes of a particular configuration to within the threshold. Thus, the within can ensure that it does not maintain too many idle processes that are unlikely to be allocated.

[00107] It should be appreciated that the method 700 shown in Figure 7 may be executed concurrently for each of the different types of available translation and transcription processes 414, 416 provided by the virtual conference provider 310 to ensure that a suitable number of such processes are available at a particular time, but that processing and memory resources are not being wasted.

[00108] Referring now to Figure 8, Figure 8 shows an example computing device 800 suitable for use in example systems or methods for providing real-time translation during virtual conferences according to this disclosure. The example computing device 800 includes a processor 810 which is in communication with the memory 820 and other components of the computing device 800 using one or more communications buses 802. The processor 810 is configured to execute processor-executable instructions stored in the memory 820 to perform one or more methods for providing real-time translation during virtual conferences according to different examples, such as part or all of the example methods 800, 900 described above with

respect to Figures 6 and 7. The computing device 800, in this example, also includes one or more user input devices 850, such as a keyboard, mouse, touchscreen, microphone, etc., to accept user input. The computing device 800 also includes a display 840 to provide visual output to a user.

[00109] In addition, the computing device 800 includes a virtual conferencing application 860 to enable a user to join and participate in one or more virtual spaces or in one or more conferences, such as a conventional conference or webinar, by receiving multimedia streams from a virtual conference provider, sending multimedia streams to the virtual conference provider, joining and leaving breakout rooms, creating video conference expos, etc., such as described throughout this disclosure, etc.

[00110] The computing device 800 also includes a communications interface 840. In some examples, the communications interface 830 may enable communications using one or more networks, including a local area network ("LAN"); wide area network ("WAN"), such as the Internet; metropolitan area network ("MAN"); point-to-point or peer-to-peer connection; etc. Communication with other devices may be accomplished using any suitable networking protocol. For example, one suitable networking protocol may include the Internet Protocol ("IP"), Transmission Control Protocol ("TCP"), User Datagram Protocol ("UDP"), or combinations thereof, such as TCP/IP or UDP/IP.

[00111] While some examples of methods and systems herein are described in terms of software executing on various machines, the methods and systems may also be implemented as specifically-configured hardware, such as field-programmable gate array (FPGA) specifically to execute the various methods according to this disclosure. For example, examples can be implemented in digital electronic circuitry, or in computer hardware, firmware, software, or in a combination thereof. In one example, a device may include a processor or processors. The processor comprises a computer-readable medium, such as a random access memory (RAM) coupled to the processor. The processor executes computer-executable program instructions stored in memory, such as executing one or more computer programs. Such processors may comprise a microprocessor, a digital signal processor (DSP), an application-specific integrated circuit (ASIC), field programmable gate arrays

(FPGAs), and state machines. Such processors may further comprise programmable electronic devices such as PLCs, programmable interrupt controllers (PICs), programmable logic devices (PLDs), programmable read-only memories (PROMs), electronically programmable read-only memories (EPROMs or EEPROMs), or other similar devices.

[00112] Such processors may comprise, or may be in communication with, media, for example one or more non-transitory computer-readable media, that may store processor-executable instructions that, when executed by the processor, can cause the processor to perform methods according to this disclosure as carried out, or assisted, by a processor. Examples of non-transitory computer-readable medium may include, but are not limited to, an electronic, optical, magnetic, or other storage device capable of providing a processor, such as the processor in a web server, with processor-executable instructions. Other examples of non-transitory computer-readable media include, but are not limited to, a floppy disk, CD-ROM, magnetic disk, memory chip, ROM, RAM, ASIC, configured processor, all optical media, all magnetic tape or other magnetic media, or any other medium from which a computer processor can read. The processor, and the processing, described may be in one or more structures, and may be dispersed through one or more structures. The processor may comprise code to carry out methods (or parts of methods) according to this disclosure.

[00113] The foregoing description of some examples has been presented only for the purpose of illustration and description and is not intended to be exhaustive or to limit the disclosure to the precise forms disclosed. Numerous modifications and adaptations thereof will be apparent to those skilled in the art without departing from the spirit and scope of the disclosure.

[00114] Reference herein to an example or implementation means that a particular feature, structure, operation, or other characteristic described in connection with the example may be included in at least one implementation of the disclosure. The disclosure is not restricted to the particular examples or implementations described as such. The appearance of the phrases “in one example,” “in an example,” “in one implementation,” or “in an implementation,” or variations of the same in various places in the specification does not necessarily

refer to the same example or implementation. Any particular feature, structure, operation, or other characteristic described in this specification in relation to one example or implementation may be combined with other features, structures, operations, or other characteristics described in respect of any other example or implementation.

[00115] Use herein of the word "or" is intended to cover inclusive and exclusive OR conditions. In other words, A or B or C includes any or all of the following alternative combinations as appropriate for a particular usage: A alone; B alone; C alone; A and B only; A and C only; B and C only; and A and B and C.

CLAIMS

That which is claimed is:

1. A method comprising:

executing, by a conference provider, a plurality of transcription and translation processes, each executed transcription and translation process associated with one of a plurality of input languages, each translation process associated with one of a plurality of output languages;

hosting, by a conference provider, a plurality of virtual conferences, each virtual conference between a plurality of client devices exchanging audio streams;

receiving, during one or more of the virtual conferences, requests from client devices to translate audio streams within the respective conferences;

allocating, during the respective conferences and in response to the respective requests to translate, one or more transcription processes and one or more translation processes to the respective virtual conferences based on a source language being used in the respective virtual conference;

providing, during the respective virtual conferences, one or more audio streams from the respective virtual conferences to the respective allocated transcription process;

providing, during the respective virtual conferences, an output from the respective transcription process to the respective allocated translation process; and

providing, during the respective virtual conferences, an output from the respective translation processes to the respective requesting client device.

2. The method of claim 1, further comprising:

determining a number of available but unallocated transcription processes;

responsive to determining a first threshold is satisfied, instantiating a second plurality of transcription processes;

determining a number of available but unallocated transcription processes; and

responsive to determining a second threshold is satisfied, instantiating a second plurality of transcription processes.

3. The method of claim 1, further comprising:

determining a number of available but unallocated transcription processes;

responsive to determining a first threshold is satisfied, terminating a second plurality of transcription processes;

determining a number of available but unallocated transcription processes; and

responsive to determining a second threshold is satisfied, terminating a second plurality of transcription processes.

4. The method of claim 1, wherein a first virtual conference of the one or more virtual conferences is allocated a first transcription process and a first translation process in response to a first request from a first client device to translate one or more audio streams within the first conference from a first source to a first target language, further comprising:

receiving, during the first virtual conference, a second request from a second client device to translate audio streams within the respective conferences, the second request identifying a second target language;

allocating, during the first virtual conference and in response to the second request to translate, a second translation process to the first virtual conference based on the first source language and the second language;

providing, during the respective virtual conferences, the output from the first transcription process to the second translation process; and

providing, during the respective virtual conferences, an output from the second translation process to the second client device.

5. The method of claim 1, wherein a first virtual conference of the one or more virtual conferences is allocated a first transcription process and a first translation process in response to a first request from a first client device to translate one or more audio streams within the first conference from a first source language to a first target language, further comprising:

receiving, during the first virtual conference, a second request from a second client device to translate audio streams within the respective conferences, the second request identifying the target language; and

providing, during the respective virtual conferences, an output from the first translation process to the second client device.

6. The method of claim 1, wherein a first virtual conference of the one or more virtual conferences is allocated a first transcription process and a first translation process in response to a first request from a first client device to translate one or more audio streams within the first conference from a first source language to a first target language, further comprising:

receiving, during the first virtual conference, a second request from a second client device to translate audio streams within the respective conferences, the second request identifying a second source language and a second target language;

allocating, during the first virtual conference and in response to the second request to translate, a second translation process and a second translation process to the first virtual conference based on the second source language and the second language;

providing, during the respective virtual conferences, one or more audio streams from the respective virtual conferences to the second allocated transcription process;

providing, during the respective virtual conferences, the output from the second transcription process to the second translation process; and

providing, during the respective virtual conferences, an output from the second translation process to the second client device.

7. The method of claim 1, further comprising:

generating, during the respective virtual conference, a transcript comprising the output from the one or more allocated transcription processes and the output from the one or more allocated translation processes.

8. A system comprising:

one or more servers, each comprising a communications interface; a non-transitory computer-readable medium communicatively coupled to the communications interface and the non-transitory computer-readable medium, the one or more servers configured to:

execute a plurality of transcription and translation processes, each executed transcription and translation process associated with one of a plurality of input languages, each translation process associated with one of a plurality of output languages;

host a plurality of virtual conferences, each virtual conference between a plurality of client devices exchanging audio streams;

receive, during one or more of the virtual conferences, requests from client devices to translate audio streams within the respective conferences;

allocate, during the respective conferences and in response to the respective requests to translate, one or more transcription processes and one or more translation processes to the respective virtual conferences based on a source language being used in the respective virtual conference;

provide, during the respective virtual conferences, one or more audio streams from the respective virtual conferences to the respective allocated transcription process;

provide, during the respective virtual conferences, an output from the respective transcription process to the respective allocated translation process; and

provide, during the respective virtual conferences, an output from the respective translation processes to the respective requesting client device.

9. The system of claim 8, wherein the one or more servers are further configured to:

determine a number of available but unallocated transcription processes;

responsive to a determination that a first threshold is satisfied, instantiate a second plurality of transcription processes;

determine a number of available but unallocated transcription processes; and

responsive to a determination that a second threshold is satisfied, instantiate a second plurality of transcription processes.

10. The system of claim 9, wherein the number of available but unallocated transcription and translation processes are associated with a first language, and wherein the one or more servers are further configured to:

select the first and second thresholds from a plurality of thresholds based on a local time associated with a first country, the first country having the first language as a primary language.

11. The system of claim 8, wherein a first virtual conference of the one or more virtual conferences is allocated a first transcription process and a first translation process in response to a first request from a first client device to translate one or more audio streams within the first conference from a first source to a first target language, wherein the one or more servers are further configured to:

receive, during the first virtual conference, a second request from a second client device to translate audio streams within the respective conferences, the second request identifying a second target language;

allocate, during the first virtual conference and in response to the second request to translate, a second translation process to the first virtual conference based on the first source language and the second language;

provide, during the respective virtual conferences, the output from the first transcription process to the second translation process; and

provide, during the respective virtual conferences, an output from the second translation process to the second client device.

12. The system of claim 8, wherein a first virtual conference of the one or more virtual conferences is allocated a first transcription process and a first translation process in response to a first request from a first client device to translate one or more audio streams within the first conference from a first source language to a first target language, wherein the one or more servers are further configured to:

receive, during the first virtual conference, a second request from a second client device to translate audio streams within the respective conferences, the second request identifying the target language; and

provide, during the respective virtual conferences, an output from the first translation process to the second client device.

13. The system of claim 8, wherein a first virtual conference of the one or more virtual conferences is allocated a first transcription process and a first translation process in response to a first request from a first client device to translate one or

more audio streams within the first conference from a first source language to a first target language, wherein the one or more servers are further configured to:

receive, during the first virtual conference, a second request from a second client device to translate audio streams within the respective conferences, the second request identifying a second source language and a second target language;

allocate, during the first virtual conference and in response to the second request to translate, a second translation process and a second translation process to the first virtual conference based on the second source language and the second language;

provide, during the respective virtual conferences, one or more audio streams from the respective virtual conferences to the second allocated transcription process;

provide, during the respective virtual conferences, the output from the second transcription process to the second translation process; and

provide, during the respective virtual conferences, an output from the second translation process to the second client device.

14. A non-transitory computer-readable medium comprising processor-executable instructions configured to cause one or more processors to:

execute a plurality of transcription and translation processes, each executed transcription and translation process associated with one of a plurality of input languages, each translation process associated with one of a plurality of output languages;

host a plurality of virtual conferences, each virtual conference between a plurality of client devices exchanging audio streams;

receive, during one or more of the virtual conferences, requests from client devices to translate audio streams within the respective conferences;

allocate, during the respective conferences and in response to the respective requests to translate, one or more transcription processes and one or more translation processes to the respective virtual conferences based on a source language being used in the respective virtual conference;

provide, during the respective virtual conferences, one or more audio streams from the respective virtual conferences to the respective allocated transcription process;

provide, during the respective virtual conferences, an output from the respective transcription process to the respective allocated translation process; and

provide, during the respective virtual conferences, an output from the respective translation processes to the respective requesting client device.

15. The non-transitory computer-readable medium of claim 14, further comprising processor-executable instructions configured to cause one or more processors to:

determine a number of available but unallocated transcription processes;

responsive to a determination that a first threshold is satisfied, instantiate a second plurality of transcription processes;

determine a number of available but unallocated transcription processes; and

responsive to a determination that a second threshold is satisfied, instantiate a second plurality of transcription processes.

16. The non-transitory computer-readable medium of claim 15, wherein the number of available but unallocated transcription and translation processes are associated with a first language, and further comprising processor-executable instructions configured to cause one or more processors to:

select the first and second thresholds from a plurality of thresholds based on a local time associated with a first country, the first country having the first language as a primary language.

17. The non-transitory computer-readable medium of claim 14, further comprising processor-executable instructions configured to cause one or more processors to:

determine a number of available but unallocated transcription processes;

responsive to a determination that a first threshold is satisfied, terminating a second plurality of transcription processes;

determine a number of available but unallocated transcription processes; and

responsive to determination that a second threshold is satisfied, terminate a second plurality of transcription processes.

18. The non-transitory computer-readable medium of claim 14, wherein a first virtual conference of the one or more virtual conferences is allocated a first transcription process and a first translation process in response to a first request from a first client device to translate one or more audio streams within the first conference from a first source to a first target language, further comprising processor-executable instructions configured to cause one or more processors to:

receive, during the first virtual conference, a second request from a second client device to translate audio streams within the respective conferences, the second request identifying a second target language;

allocate, during the first virtual conference and in response to the second request to translate, a second translation process to the first virtual conference based on the first source language and the second language;

provide, during the respective virtual conferences, the output from the first transcription process to the second translation process; and

provide, during the respective virtual conferences, an output from the second translation process to the second client device.

19. The non-transitory computer-readable medium of claim 14, wherein a first virtual conference of the one or more virtual conferences is allocated a first transcription process and a first translation process in response to a first request from a first client device to translate one or more audio streams within the first conference from a first source language to a first target language, further comprising processor-executable instructions configured to cause one or more processors to:

receive, during the first virtual conference, a second request from a second client device to translate audio streams within the respective conferences, the second request identifying a second source language and a second target language;

allocate, during the first virtual conference and in response to the second request to translate, a second translation process and a second translation process to the first virtual conference based on the second source language and the second language;

provide, during the respective virtual conferences, one or more audio streams from the respective virtual conferences to the second allocated transcription process;

provide, during the respective virtual conferences, the output from the second transcription process to the second translation process; and

provide, during the respective virtual conferences, an output from the second translation process to the second client device.

20. The non-transitory computer-readable medium of claim 14, further comprising:

generate, during the respective virtual conference, a transcript comprising the output from the one or more allocated transcription processes and the output from the one or more allocated translation processes.