



US 20110074915A1


(19) **United States**(12) **Patent Application Publication**
FERREN et al.(10) **Pub. No.: US 2011/0074915 A1**(43) **Pub. Date: Mar. 31, 2011**(54) **APPARATUS AND METHOD FOR
PRESENTING AUDIO IN A VIDEO
TELECONFERENCE****Publication Classification**(51) **Int. Cl.**
H04N 7/14 (2006.01)(52) **U.S. CL.** **348/14.11; 348/E07.083**(57) **ABSTRACT**

An advanced video teleconferencing (AVTC) system uniquely combines a number of features to promote a realistic "same room" experience for meeting participants. These features include an autodirector to select audio and video sources and to compose shots, a collaboration interface for each participant to communicate nonverbal information, directional LEDs to privately alert participants, audio reflected from the main display, and a collaborative table to share a view of objects or papers on a table. When implemented with sufficient bandwidth for take advantage of these features and to keep latency time low, this AVTC system results in a highly realistic and productive teleconferencing experience.

(76) Inventors: **BRAN FERREN**, Beverly Hills,
CA (US); **W. Daniel Hillis**, Encino,
CA (US)(21) Appl. No.: **12/959,121**(22) Filed: **Dec. 2, 2010****Related U.S. Application Data**

(63) Continuation of application No. 11/463,578, filed on Aug. 9, 2006, now Pat. No. 7,847,814, which is a continuation of application No. 10/911,079, filed on Aug. 3, 2004, now Pat. No. 7,092,002.

(60) Provisional application No. 60/504,085, filed on Sep. 19, 2003.






FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5

APPARATUS AND METHOD FOR PRESENTING AUDIO IN A VIDEO TELECONFERENCE

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This application is a continuation of U.S. patent application Ser. No. 11/463,578, filed Aug. 9, 2006, which is a continuation of U.S. patent application Ser. No. 10/911,079, now U.S. Pat. No. 7,092,002, filed Aug. 3, 2004, which claims benefit of U.S. provisional patent application Ser. No. 60/504,085, filed Sep. 19, 2003, all of which applications are incorporated herein in their entirety by this reference thereto.

BACKGROUND

[0002] 1. Technical Field

[0003] The invention relates to teleconferencing. More particularly, the invention relates to methods for better facilitating collaboration in teleconferences.

[0004] 2. Description of the Prior Art

[0005] Current video conferencing systems are deficient in their ability to support effective collaboration among conference participants, engendering a feeling of remoteness amongst the participants that is contrary to the intended benefits of the system. This deficiency can be attributed to a number of shortcomings that collectively diminish the effectiveness of the communication.

[0006] Specifically, current systems do not: Provide sufficiently low round trip latencies; Provide accurate audio cueing, i.e. remote participant voices often emanate from a location other than where the image of the remote participant is displayed; Allow participants to discuss and naturally interact with physical objects present at only one location; React to cues provided within speech patterns and content; Adapt the behavior of the system in response to frequently encountered meeting dynamics; Meaningfully track or consider personal information about the participants; or Offer a natural sense of eye contact between participants.

[0007] In conferences involving larger numbers of participants, several other deficiencies become apparent. Most notably, current systems do not: Allow participants to receive personalized information without disturbing other conference participants; Provide a natural means for participants to track who is present within the conference; or Provide a sense of the objects and people present and events transpiring in the facility surrounding the participants.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] FIG. 1 depicts an overview of a VTC installation with five stations, according to the invention;

[0009] FIG. 2 depicts the reflection of audio from the main display towards the participants, according to the invention;

[0010] FIG. 3 depicts an overhead view of the portion of the conference table onto which the collaborative table is projected, according to the invention;

[0011] FIG. 4 depicts the note passing software which is part of the collaboration interface, according to the invention; and

[0012] FIG. 5 depicts the shared application feature which is part of the collaboration interface, according to the invention.

SUMMARY

[0013] The invention is an advanced video teleconferencing system that allows an engaging and realistic video conferencing experience. The invention provides key visual, audio, security, storage, and administrative services needed for a video conferencing experience that cannot be attained with today's conventional methods. These include: an autodirector that automatically selects videocamera shots based on audio activity and other criteria, which are adjustable with user-settable operating modes; an interface to allow each participant to communicate nonverbally and privately with any other participants; directionally accurate audio reflected from the main display; a collaborative table to allow interaction with 2-D and 3-D objects; and directional LEDs to privately alert participants. This functionality is provided through interfaces that are simple, intuitive, and easily operated by those with little or no computer literacy.

DETAILED DESCRIPTION

[0014] The herein disclosed advanced video teleconferencing (AVTC) system provides a same room experience for meeting participants at physically separated locations, thereby greatly enhancing user interaction and collaboration. This is met through the combination of a number of features, each of which is described in detail below: Low latency to enhance the immediacy of the videoconference experience. An autodirector to automatically select audio and video sources and to compose shots, based on a number of criteria. These criteria include audio activity, operating modes, participant identification, and other factors. Participant transponders or login to incorporate personal information into the autodirector criteria, including relative importance of the participant and security clearance. A collaboration interface for participants to communicate nonverbal information to any other participants. Directional LED notification to privately alert participants. Audio reflected from the main display, such that audio appears to originate from the speakers on the screen. Image-based participant identification for easy reference by participants. A collaborative table to share a view of objects or papers on a table, along with hand gestures.

[0015] For the purposes of this description, a "site" is a single physical location at which a VTC system is situated and a "station" is a location at a site designed to accommodate a single participant with video and audio pickups and other tools. Each site has a minimum of one station.

[0016] FIG. 1 depicts an overview of a VTC site with five stations, according to the invention. Each station **10** accommodates one participant, and has a collaboration interface **11**. In the preferred embodiment, the interface is a touchscreen via which the participant can send and receive notes, annotate files and images, and perform other tasks. The stations are arranged around a semicircular table **12**, which helps facilitates the illusion that participants from other sites are seated around the same table. Other sites and participants are shown on the main display **13**. Supplemental displays **14** can show additional information, such as maps, stills, or moving images.

[0017] Narrow field videocameras **15** are aimed at each participant. A directional LED is mounted atop each video-

camera to privately alert the corresponding participants. Wide field videocameras **16** capture the room context and can be panned or tilted. Audio is received through microphones for each participant, as well as room microphones to pick up conversation from observers not seated at stations and other ambient sounds. A speaker **17** below the surface of the table is aimed at the main display, thus reflecting audio back to the participants to give the illusion that the audio originates from the display itself.

[0018] The collaborative table display **18** is projected on the table **12** from an overhead projector (not shown), or alternatively from a projector beneath the table. The collaborative table is used to show maps and other flat images projected to or from another site, and shows any objects held over the area, as well as the hands of the participant pointing out features or manipulating objects.

Adequate Bandwidth for Reduced Latency

[0019] Traditionally, there is a perceived tradeoff in video conferencing systems between available bandwidth and video quality. Given an available bandwidth, the image resolution and frame rate are selected such that an aggressive, efficient compression scheme can transmit the video signal within that bandwidth. However, experience has shown that low bandwidth systems with aggressive and efficient compression and decompression (CODEC) schemes still yield poor video conferencing experiences.

[0020] The AVTC system is predicated in part on the belief that latency is the primary cause for this shortcoming. Accordingly, the AVTC system considers a tradeoff between available bandwidth and latency. In the preferred embodiment, adequate bandwidth is provided to yield round trip latencies well underneath 100 ms, preferably as low as 30 ms. Because a large fraction of the round trip latency is attributable to the operation of the CODECs, this essentially requires bandwidth great enough that the desired image size and frame rate can be transmitted at the compression ratio achievable by the CODECs in the desired latency.

[0021] If such bandwidth is not available, the latency may be preserved by reducing the frame rate or image resolution. Preferably, this tradeoff is adjusted dynamically by the AVTC system. Alternatively it may be adjusted manually, preferably by a meeting administrator seated at an administrator's console. The ability to adjust this tradeoff may be particularly useful in network configurations, where the bandwidth used by the AVTC system is a shared resource.

Autodirector

[0022] The AVTC system incorporates an autodirector that emulates the behavior of a highly experienced human video director. The autodirector automatically selects, from among one or more videocamera feeds and other video inputs, a video signal for transmission to remote video conferencing sites. In this manner, the flow of conversation among the participants of a local video conferencing site is seamlessly reconstructed at the remote teleconferencing site for viewing by the remote participants.

[0023] The autodirector selection algorithm has been described in detail in a related U.S. patent, "Method and Apparatus for Selection of Signals in a Teleconference", patent Ser. No. 10/327,368, filed Dec. 20, 2002 and allowed

Jun. 12, 2004. The description below is focused on two other unique features of the autodirector, the operating modes and speech recognition.

Operating Modes

[0024] Because the social dynamics of any conference can change as the meeting progresses, the autodirector incorporates a sense of meeting dynamics in determining its selection of video inputs. Specific modes can be selected to match particular situations, which will adjust video and audio selection criteria used by the autodirector. In the case where a participant wishes to address and be viewed by all participants of a meeting, the autodirector is configured to force all meeting participants to view the participant making the address. This mode is termed speech mode. A related mode is termed lecture mode, and handles situations where a participant wishes to address another single participant to make an important point. The system is thereby configured for a one-on-one exchange between participants. The two participants in the exchange see each other, and participants at other sites see both individuals alternately, according to the dynamics of the verbal exchange. Another mode offered by the autodirector is view force mode, where an administrator at one site can force the selection of a particular shot from a specified site. The autodirector continues to normally switch between sites, but whenever the specified site is selected, all other sites view only the particular shot from that site. Finally, in introductory fast switching mode, a rapid switching algorithm is provided that allows the system to track the rapid exchanges that typically occur in the initial phase of a conference, in which individuals introduce themselves and meet one another.

[0025] The autodirector can be placed in any of these modes by a participant located at any site, via an administrator console. Priority is given to the latest received mode request. The current mode is displayed on the administrator console at each site. An administrator can create new modes or modify existing modes.

[0026] If the administrator does not specify a mode, the autodirector preferably selects the appropriate mode based on assumed or detected meeting dynamics. For example, in a typical business meeting, conversation often begins with a series of introductions among participants, often as a sequence of rapidly changing speakers. To capture these events effectively, the audio and video signal selected should also be changed rapidly. The autodirector may therefore automatically enter introductory fast switching mode as a meeting begins. Thus, the autodirector, aware of introductions as part of typical meeting dynamics, initially allows signal switching patterns that might appear abrupt or hurried under other circumstances.

[0027] Later in a typical meeting, meeting dialogue becomes more focused and directed. The duration for which an individual speaks typically increases as more complex issues are addressed, and extended presentations may be made. The autodirector may therefore readily enter speech mode and lecture mode during the middle of a typical meeting. An appropriate time to enter into these modes may be detected by analyzing speech patterns.

[0028] The AVTC system also stores specialized sequences of modes for use with a particular style of meeting. For example, one sequence of modes may be selected for use during a short staff meeting, while a different sequence of modes may be used for an extended board of directors meeting. These sequences are fully customizable and are stored

within the system for later access. The system can therefore be made aware of dynamics specific to a particular group of participants.

[0029] The resulting system is an automated video conferencing direction system that functions in accord with, rather than in spite of, typical meeting dynamics.

Speech Recognition

[0030] To provide an enhanced understanding of which video signal is appropriate for selection at a particular time, the autodirector analyzes the conference audio. In one embodiment, the autodirector favors a shot of a participant when his or her name is detected on the audio; this is effective at capturing reaction shots. Limiting speech recognition processing to searching the audio for participant names greatly reduces the complexity of this problem.

[0031] More generally, specific words or phrases may be recognized and treated as attractors or pointers. Attractors increase the likelihood of selecting a video signal containing an image of the participant who uttered the attractor word or phrase, such as “I think,” “In my opinion,” or “I have here”. Pointers, of which participants names are one example, encourage the selection of a particular participant other than the participant who uttered the phrase.

Audio Reflected from Monitor for Improved Audio Directionality

[0032] When audio does not seem to be coming directly from the same direction as the main display, such as from side speakers, fatigue can be generated among the participants. Conversely, proper audio directionality results in more natural, effortless communication between participants.

[0033] To provide accurate audio directionality, the AVTC system incorporates audio speakers directed at the main display showing images of the remote participants. Audio emanating from the speakers is reflected from the display and towards the listening participants, thus appearing to emanate directly from the participant shown on the main display. This effect is further enhanced by selecting a highly directional audio speaker and shielding the speaker in a manner that blocks the straight line transmission pathways between the speaker and the listening participants.

[0034] FIG. 2 depicts the reflection of audio from the main display towards the participants according to the invention. The speaker 21 is mounted beneath the surface of the table under an acoustically transparent material 22 and aimed at the main display 12. Audio is thereby reflected from the main display towards the participants 23, giving the illusion that the audio is originating from the center of the display.

Participant Transponders

[0035] To obtain and meaningfully incorporate participant personal information into the conference environment, one embodiment of the AVTC system uses participant transponders. Each transponder provides information about a video conferencing participant, such as his location within the conference site, organizational position and title, clearance level, and speaking characteristics. The information is provided to the autodirector, which can then make more intelligent decisions based on the personal information.

[0036] The participant transponder is worn or carried by the video conferencing participant, for example in the form of a radio frequency identification (RFID) tag embedded in a personnel ID card. More generally, the device may be incorpo-

rated within a badge, pen, laser pointer, wireless microphone, or other device common to a video conference setting. In the preferred configuration of the AVTC system, the transponder is integrated within a station at the conferencing site through the use of existing RFID card technology, either passive or active. In this scenario, the participant carries an RFID card, and the card reader is incorporated within the station, which is in turn connected to the AVTC system.

[0037] Regardless of the form factor or communication mechanism, the transponder need not broadcast all information about the participant, but instead can broadcast a participant identity that allows the AVTC system to access a complete participant profile within a database. From this database the autodirector can determine a participant hierarchy, as well as a security level for the conference as a whole. The conference security level is displayed for all participants by the autodirector. The autodirector can accept changes to the hierarchy, allowing participants to take on varying priority levels as a meeting evolves.

[0038] The database may also contain biometric data for the participants. This allows the identity of the participant to be verified by biometric measurements, such as facial features (e.g. as captured by one or more of the video cameras), or fingerprint or retinal scans.

[0039] The transponder may also allow the location of the wearer within the conference room to be determined. This can be performed locally by the transponder, and then broadcast to the AVTC system when queried. Alternatively, the location can be determined through time of flight or directional triangulation calculations based on measurements acquired from sensors placed throughout the conferencing site and operated by the AVTC system. The transponders may also be used to detect the entry or exit of a participant by analyzing the positions of the transponder signals it receives. In particular, the security level of the meeting is adjusted to reflect the participants currently present. Based on the current security level, the AVTC system allows or prohibit access to specific information or services.

Collaboration Interface

[0040] Each station is fitted with a collaboration interface for sending nonverbal information to other participants. The interface consists of a pen-enabled touchscreen display, preferably to one side of the station so as not to block the view of the participant. The interface provides access to a suite of collaboration tools: Login—If participant transponders are not used, the participants can use this screen to login or can run an identification badge through a card reader at the side of the interface. The system can then display the person's name and title to other participants. Personalized services and preferences can be set for specific participants, such as directing their important email to their interface. When the participant logs out from the conference, the autodirector can cease selecting the view of that station. Alternatively, a presence indicator can detect if a person is sitting at the station. Note sending—Each participant can use a virtual whiteboard to send notes or hand-drawn illustrations to any or all participants, who are alerted by a directional LED mounted above each station's videocamera. The contents of the whiteboard can be printed on a local printer, or cleared by a “Clear” button.

[0041] FIG. 4 depicts the note passing software which is part of the collaboration interface, according to the invention. The note passing mode is selected from tabs 40 at the top of

the interface. Notes or sketches can be handwritten on the notes area **41** using an electronic pen tool or similar device (not shown). To send a note, the participant selects another participant from the image-based identification area **42**, which contains thumbnail images of each participant. The note can then be sent or printed through control buttons **43**. Note recipients are privately alerted by the flashing of a directional LED aimed at their station. Web browser—Any participant can view, annotate, or send Web pages to other participants. Shared applications—Participants can operate software applications, such as Microsoft PowerPoint presentations, on their interface with the same annotation and distribution capabilities described above. FIG. 5 depicts the shared application feature which is part of the collaboration interface, according to the invention. In this depiction, the shared application is slide presentation software. Participants can send any image or slide **50** to the collaboration screen. All participants see the contents of the collaboration screen, which is akin to an electronic whiteboard. Participants can handwrite annotations **51**, which appear in different color inks to distinguish participant's annotations.

Image-Based Participant Identification

[0042] To provide identification of the participants within a conference, the AVTC uses an image-based approach. A still image of a participant is captured using one of the several video cameras present. Preferably, the image is captured by the camera positioned to capture a close up shot of the participant. The image is then reduced in size and converted to a suitable format for presentation on a computer display, for example TIFF, GIF, or JPEG. The image may then be presented on the displays of other users as an iconic representation of the presence of the participant in the meeting. For example, the image may be presented at the remote location among a set of pictures indicating who is present at the local location. Participants can indicate the recipients of items such as emails and notes by selecting the image of other participants. FIG. 4 depicts an embodiment of participant images **42**.

[0043] Alternatively, images of the participants may be retrieved from a database of user information, based on a username obtained from participant transponders or by the login or ID card reader mechanisms described herein. Alternatively, the image may be stored directly on the ID card.

[0044] Regardless of the technique used to obtain the image of the participant, by presenting images in addition to names, the AVTC system provides a natural way for one participant to determine other participants currently present. Because many individuals find faces easier to recognize and remember than names, this approach allows participants to interact with one another in a more comfortable manner.

Directional LED Participant Notification

[0045] During video conferences with multiple participants at a single site, it is often desirable to obtain the attention of a single participant without disturbing the other participants. For example, the AVTC system may wish to notify a participant that he is needed in another meeting, that new email has arrived, or that a note has been received via the collaborative interface.

[0046] To provide such notification, the AVTC system incorporates a directional LED-based cueing device. An LED is positioned at the end of a relatively long tube, with the axis

of LED illumination directed along the length of the tube. The axis of the tube is aligned with the participant to be alerted so that when the LED is illuminated, it is visible only to the intended participant. The length of the tube is selected based on the spacing between participants and the distance from the tube to the participants.

[0047] Preferably, the system provides one notification device for each participant. Alternatively, a single tube and LED may be actuated, such that it may be aligned with a particular participant to be notified. Notably, the preferred embodiment offers the advantage of simultaneous notification of more than one participant. Further, the LED and tube are preferably positioned near the main display, so that the notification device is usually within the field of view of the participant. In one embodiment, the tubes are mounted along the optical axis of the cameras aimed at each participant. The notification device can then further function as an alignment device, wherein each participant may adjust his position until he or she is aligned with the tube, ensuring a properly composed individual shot.

[0048] Alternatively, an alert may be presented on a personal display located alongside the participant. However, because the primary focus of the participant is on the main display, LED-based notification is preferred, because it more likely to be immediately noted by the participant.

Camera Coverage of Secondary Participants

[0049] Since wide views do not give sufficient participant details, most shots in a videoconference are close-ups. However, these do not give remote participants a feel for the surrounding conference room environment and of secondary participants.

[0050] To address this problem, the AVTC includes one or more room microphones and cameras that provide audio and camera coverage of secondary participants not positioned at stations. For example, if the primary participants are seated at a conference table centered in front of a monitor positioned flush with one wall of a conference room, the room cameras provide coverage of the corners of the room on the side opposite the monitor.

[0051] The video and audio signals provided by the room cameras and microphones are provided to the autodirector. The autodirector then determines appropriate times to transmit these video and audio feeds to the remote location, providing a sense of context to the remote participants. For example, the room camera feed may be selected when the primary participants are quiet, when a secondary participant speaks for an extended period of time, or if the entry or exit of a participant through a conference room doorway is detected.

Collaborative Table

[0052] The AVTC provides participants with a collaborative table upon which objects and hand gestures can be viewed at several sites.

[0053] For convenience, the table at one site is termed the source table, and the table at a second site is termed the viewing table. A participant sits at the source table and places objects or documents of interest upon it. Above the source table is a high definition table videocamera which captures images of the objects, which are encoded and transmitted. At the other site or sites, a projector, preferably high-resolution and overhead mounted, projects the images onto the viewing table.

[0054] FIG. 3 depicts an overhead view of the portion of the conference table onto which the collaborative table is projected, according to the invention. The image 18 is projected on the conference table 12 from a projector overhead from another video teleconference site, and is centered in front of the participant 20 situated at the middle of the conference table. The high definition image shows both still objects in the background 21 and the hands of a participant at the other site manipulating a solid object 22.

[0055] If the source participant gestures with his hands over the source table, then the participant at the viewing table sees the source participant's hands. Given limited bandwidth, this leads to a seemingly conflicting requirements whereas objects of interest should be presented at the maximum possible quality (low frame rate), yet the hand gestures should appear smooth and fluid (high frame rate). In addition, the participant at the viewing table should in turn be able to make hand gestures at the projected objects that are in turn sent and projected onto the source table.

Encoding

[0056] Described below are three embodiments of an encoding scheme for the collaborative table which provide for both high resolution, largely static scenes, with lower resolution but high frame rate hand gestures and motion. These approaches range from an computationally less expensive but lower quality method (motion detection) to a high quality custom mechanism.

Motion Detection with High and Low Frame Rate Encodings

[0057] Open source MJPEG code is used to generate a motion JPEG RTP stream with tight custom encoding. MJPEG allows different image resolutions in each frame, and each frame is JPEG encoded with no intraframe dependencies. The system constantly captures images and detects significant changes between frames. If no significant change has occurred over a period of time, then the system enters static mode, and sends a high quality JPEG frame onto the channel via RTP. A single high quality frame takes several seconds to send at T1 bandwidths, and is sent about every 30 seconds so that any new listeners are updated in a timely manner.

[0058] If the system detects a significant change in the frame above a certain threshold, it immediately switches to dynamic mode, encoding high frame rate, lower resolution images into the RTP MJPEG stream. Listening clients automatically adjust to the new frame resolution. Because the system is MJPEG compliant, it can use an existing RTP/MJPEG client with little or no modification. If new motion is not detected after several seconds, the system switches back to static mode and starts sending high quality, very low rate frames.

[0059] A disadvantage of this technique is that if, for example, the source table consists of a map over which someone is gesturing, when the person gestures over part of the image with their hands the entire image, including the map, is degraded to illegible quality until the person removes their hand or holds their hand very still.

Masked Hybrid: High and Low Frame Rate Streams with Masking

[0060] In dynamic mode, this technique uses a new RTP channel called the mask channel, containing low resolution run-length-encoded bit-masks to correspond to each of the transmitted high frame-rate, low quality MJPEG dynamic frames. When the source system enters dynamic mode, it retains the last transmitted static high quality frame as a

reference image and compares all successive dynamic frames to a scaled down reference frame. It computes and thresholds the pixel-by-pixel change between the reference frame, scaled to the dynamic mode resolution, and each dynamic frame and builds a bit-mask of changed areas. This bit-mask is labeled and transmitted on the mask channel after the dynamic frame is sent. Each RTP display client receives the mask and the dynamic frame, and correlates the two according to the label. The client then copies the high resolution reference frame, and overlays the scaled-up dynamic frame according to the mask.

[0061] This technique allows stationary objects on the source table to retain their full-resolution appearance while low fidelity, high frame rate hand gestures are overlaid in a lower resolution.

Custom Encoding

[0062] A higher quality solution is obtained with an entirely custom coding strategy. The high resolution scene is decomposed into a multiscale image, composed of a stack of images with the lowest quality base image overlaid by successive images. These images contain higher order terms refining the quality of the base image up to any level of desired quality, very similar to a wavelet deconstruction.

[0063] The base image is small enough to allow a full frame, high rate transmission in any circumstance, e.g. 135.times.238.times.2 bit greyscale. Higher order images successively increase depth and resolution in scale space up to full high-definition, 1080i resolution. For a transmission model, the images are subdivided into about 1024 spatial blocks.

[0064] To create the multiscale image, starting at the base layer the encoder computes a difference value from the prior frame for each spatial block. If the difference value for a block at that layer is significant, then the encoder flags the block and all corresponding higher layer image blocks as needing to be retransmitted with a certain disparity score. The encoder then starts at the base layer and transmits via RTP each changed block in that layer in the order of it's disparity score and marks the transmitted blocks as current.

[0065] If there is time before the next frame, the encoder then goes to the second order image in the stack and transmits any flagged blocks until it either runs out of bandwidth before the next frame, or runs out of flagged blocks on that layer image. If there is time before the next frame, then the encoder traverses higher and higher layer images and transmits changed blocks, dynamically capturing more and more image detail.

[0066] If the scene becomes static, this transmission algorithm automatically builds and transmits a high quality representation of the image over several seconds because small changes such as lighting and camera noise, should not affect the lower layer base images. This gives the encoder time to transmit blocks from the detailed higher order images.

[0067] In the case of a largely static scene with a hand moving over it, performance is boosted by comparing each block not just with the prior frame, but also with a reference frame that has been built up over time. If a hand is removed from over a portion of a static scene, and that portion of the image corresponds well with the reference image, then those blocks can be flagged to revert to their detailed reference state instead of being queued for full retransmission. Each client maintains the shared multiscale representation, adjusting it as new blocks or "revert block to reference state" flags are sent.

[0068] This algorithm provides optimal performance for the stated requirements, providing high frame rate when change is prevalent, but at the maximum allowable quality given the amount of change in the scene, while preserving a high quality reference state.

Simultaneous Projection and Image Capture

[0069] Challenges arise when simultaneously capturing and projecting video onto the collaborative table. For example, a map on the source table is projected onto the viewing table, while simultaneously the hands of the participant at the viewing table (pointing at the projected map) are projected onto the source table. The problem is that not only the map image, but also the hands of the participant at the viewing table projected onto the source table, are transmitted back to the viewing table.

[0070] In theory, both the actual map on the source table and the projected hands could be captured, subtract the projected hands, then transmit the map image without the hands. In reality, this is a nontrivial if not incomplete problem, especially when the projection include three-dimensional items of unknown shape, varying reflectivity, and unknown color.

[0071] Instead, the system momentarily blanks the projector to black when the camera is capturing a frame, and runs both the table videocamera and the projector at a reduced duty cycle and frame rate. By synchronizing the camera capture, only 30 ms are needed to capture a frame. Because encoding and transmitting at full frame rate may be impractical for the reasons described earlier, capture is preferably executed at 5 frames/sec by blanking the projector 5 times a second for 30 ms each time. This blanking is generally imperceptible to the viewer of the projection.

Projector Blanking Issues

[0072] Implementing the projector blanking solution presents two practical problems: Rapid and precise blanking—Successfully blanking the projector by interrupting the video to it for a period of a video frame requires a rapid response from the projector. While LCD projectors typically have somewhat slow response times, a broadcast frame is considerably longer than a typical high resolution graphics frame, so an LCD projector may be suitable. However, using a higher quality CRT-based projector is preferable due to their faster response times.

[0073] For the software driving the projector, swapping the video output to a blank buffer for a few frames and outputting a synchronization trigger signal for the camera can be readily achieved with contemporary consumer graphics cards. If the image buffer is to be redrawn while the blank buffer is being displayed, the double buffering functionality supported by the majority of graphics cards may be used.

[0074] Another approach to blanking the projector without concern for the projector response time incorporates an LCD shutter over the projector lens. This reduces the gain of the projection, but existing stereo vision LCD shutter glasses demonstrate that LCD shutter response time is far more than adequate for a broadcast frame blanking. Synchronized blanking—Synchronization of the blanking with image capture is achieved using an external sync signal. However the sync signal is generated, e.g. by the camera or by an external source, the signal is input into a counter to signal the computer to blank the projection periodically and capture the next frame, with suitable projector refresh timing. Latencies in the

capture hardware may complicate this process slightly. Because the horizontal refresh rate for the projector is likely to be two or three times the broadcast frame capture rate, delay due to projector frame rate should be minimal. Regardless, with a fast enough projector, the timing problem is readily addressed.

[0075] An alternate embodiment polarizes the light from the projector, and use a polarizing filter on the camera to filter out the polarized projected image, which may be preferable to the more complex projector blanking and camera synchronization solution.

Background Segmentation

[0076] At the viewing table, the system simultaneously captures the viewer's hand gesturing with the camera while projecting the source table's images onto the table. The same projector/camera synchronization described above is used to separate the hands from the projection. The background for the viewing table is presumed to be white or some other solid color. However, an additional problem is encountered at the viewing end, namely that only the viewer's hands should be projected down onto the source table, without adding a lot of light and flicker by projecting the white background onto the source table.

[0077] Thus, the viewer's hands are masked from the (presumably) white background, and replaced with a black mask. This is easily accomplished through simple thresholding. If the background is a color other than white, color segmentation and connected component labeling can be performed on the viewer side before encoding of the hand video. If the background is not a solid color, background segmentation is slightly more complicated, but may be performed with the thresholded background masking described herein in connection with encoding.

Camera-Projector Calibration

[0078] To ensure proper functioning of the hand overlays from the viewing table to the source table, it is necessary to calibrate the relative position and scale of the projected image and the camera capturing the hand motion on the viewing table. This calibration may be automated by projecting a set of patterns onto the table and detecting their position in the camera with simple image processing. Merely projecting a solid screen and detecting the corners provides most of the necessary parameters.

[0079] Although the invention is described herein with reference to several embodiments, including the preferred embodiment, one skilled in the art will readily appreciate that other applications may be substituted for those set forth herein without departing from the spirit and scope of the invention.

[0080] Accordingly, the invention should only be limited by the following claims.

1. A video teleconferencing system comprising:
 - a first display screen within view of each participant of a video conference;
 - at least one videocamera paired with said each participant of said video conference;
 - at least one microphone paired with said each participant of said video conference;
 - a video switching system that automatically selects a single incoming video signal originating from one of said at

least one videocamera for display on said first display screen, based on at least one of:

an audio activity level of at least one of said microphones; and

selection parameters specified by an active operating mode among a plurality of available operating modes.

2. The system of claim 1, wherein said active operating mode is specified by at least one of said participants.

3. The system of claim 1, wherein said active operating mode is automatically selected by said automated video switching system.

4. The system of claim 1, wherein said selection parameters can be adjusted by at least one of said participants.

5. The system of claim 1, wherein selection parameters for additional available operating modes can be specified by an administrator of said system.

6. The system of claim 1, wherein said automated video switching system stores sequences of said active operating modes, each of said sequences reflecting a particular meeting dynamic.

7. The system of claim 1, wherein said available operating modes comprise a speech mode wherein said selection parameters ensure that a video signal among said video signals that displays an image of a specified participant among said participants at a first site is displayed continuously to all participants at other sites.

8. The system of claim 1, wherein said available operating modes comprise a lecture mode wherein said selection parameters ensure that

two specified participants among said participants view only video signals among said video signals displaying images of each other, and

other participants among said participants alternately view video signals among said video signals displaying first one of said two specified participants and then a second one of said two specified participants.

9. The system of claim 1, wherein said available operating modes comprise a view force mode wherein said selection parameters ensure that said automated video switching system at a first of said at least one site selects for transmission a video signal specified by one of said participants at a second of said at least one site.

10. The system of claim 1, wherein said available operating modes comprise a fast switching mode wherein said selection parameters ensure that said automated video switching system selects from among said video signals at an increased frequency, so that said automated video switching system tracks rapid exchanges between said participants.

11. The system of claim 1, further comprising a transmitting system for transmitting said signal originating from one of said at least one videocamera to any of said video switching system and any of said first display screen within view of each participant.

* * * * *