



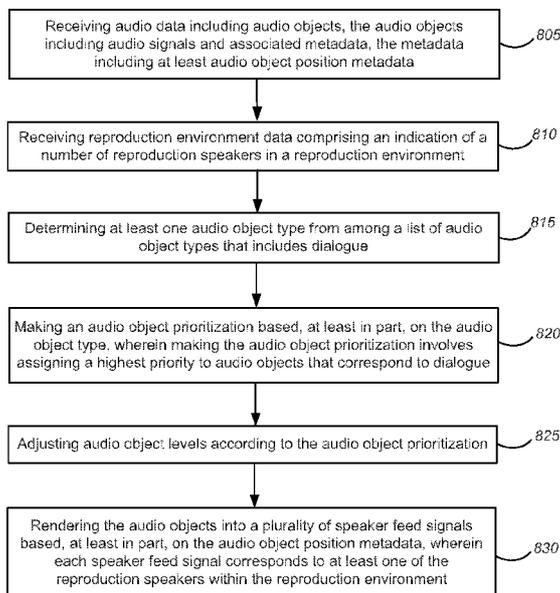
- (51) **International Patent Classification:**  
H04S 7/00 (2006.01) H04R 3/12 (2006.01)  
H04R 27/00 (2006.01)
- (21) **International Application Number:**  
PCT/US20 16/028295
- (22) **International Filing Date:**  
19 April 2016 (19.04.2016)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**  
62/149,946 20 April 2015 (20.04.2015) US
- (71) **Applicant: DOLBY LABORATORIES LICENSING CORPORATION** [US/US]; 1275 Market Street, San Francisco, California 94103 (US).
- (72) **Inventors: DE BURGH, Mark David;** c/o Dolby Australia Pty Ltd, 35 - 51 Mitchell Street, McMahon's Point Level 3, Sydney, New South Wales NSW 2060 (AU).  
**YAP, Tet Fei;** c/o Dolby Australia Pty Ltd, 35 - 51

Mitchell Street, McMahon's Point Level 3, Sydney, New South Wales NSW 2060 (AU).

- (74) **Agents: ANDERSEN, Robert L.** et al; Dolby Laboratories, INC., Intellectual Property Group, 1275 Market Street, San Francisco, California 94103 (US).
- (81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) **Designated States** (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU,

[Continued on nextpage]

(54) **Title:** PROCESSING AUDIO DATA TO COMPENSATE FOR PARTIAL HEARING LOSS OR AN ADVERSE HEARING ENVIRONMENT



800

Figure 8

(57) **Abstract:** Methods are provided for improving an audio scene for people suffering from hearing loss or for adverse hearing environments. Audio objects may be prioritized. In some implementations, audio objects that correspond to dialogue may be assigned to a highest priority level. Other implementations may involve assigning the highest priority to other types of audio objects, such as audio objects that correspond to events. During a process of dynamic range compression, higher-priority objects may be boosted more, or cut less, than lower-priority objects. Some lower-priority audio objects may fall below the threshold of human hearing, in which case the audio objects may be dropped and not rendered.

WO 2016/172111 A1



# PROCESSING AUDIO DATA TO COMPENSATE FOR PARTIAL HEARING LOSS OR AN ADVERSE HEARING ENVIRONMENT

## CROSS REFERENCE TO RELATED APPLICATIONS

5           [0001] The present invention claims the benefit of United States Provisional Patent Application No. 62/149,946, filed on April 20, 2015, which is hereby incorporated by reference in its entirety.

## TECHNICAL FIELD

10           [0002] This disclosure relates to processing audio data. In particular, this disclosure relates to processing audio data corresponding to diffuse or spatially large audio objects.

## BACKGROUND

15           [0003] Since the introduction of sound with film in 1927, there has been a steady evolution of technology used to capture the artistic intent of the motion picture sound track and to reproduce this content. In the 1970s Dolby introduced a cost-effective means of encoding and distributing mixes with 3 screen channels and a mono surround channel. Dolby brought digital sound to the cinema during the 1990s with a 5.1 channel format that provides discrete left, center and right screen channels, left and right surround arrays and a subwoofer channel for low-frequency effects. Dolby Surround 7.1, introduced in 2010, increased the number of surround channels by splitting the existing left and right surround channels into  
20 four "zones."

          [0004] Both cinema and home theater audio playback systems are becoming increasingly versatile and complex. Home theater audio playback systems are including increasing numbers of speakers. As the number of channels increases and the loudspeaker layout transitions from a planar two-dimensional (2D) array to a three-dimensional (3D) array  
25 including elevation, reproducing sounds in a playback environment is becoming an increasingly complex process.

          [0005] In addition to the foregoing issues, it can be challenging for listeners with hearing loss to hear all sounds that are reproduced during a movie, a television program, etc. Listeners who have normal hearing can experience similar difficulties in a noisy playback  
30 environment. Improved audio processing methods would be desirable.

## SUMMARY

[0006] Some audio processing methods disclosed herein may involve receiving audio data that may include a plurality of audio objects. The audio objects may include audio signals and associated audio object metadata. The audio object metadata may include audio object position metadata.

[0007] Such methods may involve receiving reproduction environment data that may include an indication of a number of reproduction speakers in a reproduction environment. The indication of the number of reproduction speakers in the reproduction environment may be express or implied. For example, the reproduction environment data may indicate that the reproduction environment comprises a Dolby Surround 5.1 configuration, a Dolby Surround 7.1 configuration, a Hamasaki 22.2 surround sound configuration, a headphone configuration, a Dolby Surround 5.1.2 configuration, a Dolby Surround 7.1.2 configuration or a Dolby Atmos configuration. In such implementations, the number of reproduction speakers in a reproduction environment may be implied.

[0008] Some such methods may involve determining at least one audio object type from among a list of audio object types that may include dialogue. In some examples, the list of audio object types also may include background music, events and/or ambiance. Such methods may involve making an audio object prioritization based, at least in part, on the audio object type. Making the audio object prioritization may involve assigning a highest priority to audio objects that correspond to dialogue. However, as noted elsewhere herein, in alternative implementations making the audio object prioritization may involve assigning a highest priority to audio objects that correspond to another audio object type. Such methods may involve adjusting audio object levels according to the audio object prioritization and rendering the audio objects into a plurality of speaker feed signals based, at least in part, on the audio object position metadata. Each speaker feed signal may correspond to at least one of the reproduction speakers within the reproduction environment. Some implementations may involve selecting at least one audio object that will not be rendered based, at least in part, on the audio object prioritization.

[0009] In some examples, the audio object metadata may include metadata indicating audio object size. Making the audio object prioritization may involve applying a function that reduces a priority of non-dialogue audio objects according to increases in audio object size.

[0010] Some such implementations may involve receiving hearing environment data that may include a model of hearing loss, may indicate a deficiency of at least one reproduction speaker and/or may correspond with current environmental noise. Adjusting the audio object levels may be based, at least in part, on the hearing environment data.

5 [0011] The reproduction environment may include an actual or a virtual acoustic space. Accordingly, in some examples, the rendering may involve rendering the audio objects to locations in a virtual acoustic space. In some such examples, the rendering may involve increasing a distance between at least some audio objects in the virtual acoustic space. For example, the virtual acoustic space may include a front area and a back area (e.g.,  
10 with reference to a virtual listener's head) and the rendering may involve increasing a distance between at least some audio objects in the front area of the virtual acoustic space. In some implementations, the rendering may involve rendering the audio objects according to a plurality of virtual speaker locations within the virtual acoustic space.

[0012] In some implementations, the audio object metadata may include audio object  
15 prioritization metadata. Adjusting the audio object levels may be based, at least in part, on the audio object prioritization metadata. In some examples, adjusting the audio object levels may involve differentially adjusting levels in frequency bands of corresponding audio signals. Some implementations may involve determining that an audio object has audio signals that include a directional component and a diffuse component and reducing a level of the diffuse  
20 component. In some implementations, adjusting the audio object levels may involve dynamic range compression.

[0013] In some examples, the audio object metadata may include audio object type  
metadata. Determining the audio object type may involve evaluating the object type  
metadata. Alternatively, or additionally, determining the audio object type may involve  
25 analyzing the audio signals of audio objects.

[0014] Some alternative methods may involve receiving audio data that may include a  
plurality of audio objects. The audio objects may include audio signals and associated audio  
object metadata. Such methods may involve extracting one or more features from the audio  
data and determining an audio object type based, at least in part, on features extracted from  
30 the audio signals. In some examples, the one or more features may include spectral flux,  
loudness, audio object size, entropy-related features, harmonicity features, spectral envelope  
features, phase features and/or temporal features. The audio object type may be selected  
from a list of audio object types that includes dialogue. The list of audio object types also

may include background music, events and/or ambiance. In some examples, determining the audio object type may involve a machine learning method.

[0015] Some such implementations may involve making an audio object prioritization based, at least in part, on the audio object type. The audio object prioritization may  
5 determine, at least in part, a gain to be applied during a process of rendering the audio objects into speaker feed signals. Making the audio object prioritization may involve assigning a highest priority to audio objects that correspond to dialogue. However, in alternative implementations making the audio object prioritization may involve assigning a highest  
10 priority to audio objects that correspond to another audio object type. Such methods may involve adding audio object prioritization metadata, based on the audio object prioritization, to the audio object metadata.

[0016] Such methods may involve determining a confidence score regarding each audio object type determination and applying a weight to each confidence score to produce a weighted confidence score. The weight may correspond to the audio object type  
15 determination. Making an audio object prioritization may be based, at least in part, on the weighted confidence score.

[0017] Some implementations may involve receiving hearing environment data that may include a model of hearing loss, adjusting audio object levels according to the audio object prioritization and the hearing environment data and rendering the audio objects into a  
20 plurality of speaker feed signals based, at least in part, on the audio object position metadata. Each speaker feed signal may correspond to at least one of the reproduction speakers within the reproduction environment.

[0018] In some examples, the audio object metadata may include audio object size metadata and the audio object position metadata may indicate locations in a virtual acoustic  
25 space. Such methods may involve receiving hearing environment data that may include a model of hearing loss, receiving indications of a plurality of virtual speaker locations within the virtual acoustic space, adjusting audio object levels according to the audio object prioritization and the hearing environment data and rendering the audio objects to the plurality of virtual speaker locations within the virtual acoustic space based, at least in part,  
30 on the audio object position metadata and the audio object size metadata.

[0019] At least some aspects of the present disclosure may be implemented via apparatus. For example, one or more devices may be capable of performing, at least in part, the methods disclosed herein. In some implementations, an apparatus may include an interface system and a control system. The interface system may include a network interface,

an interface between the control system and a memory system, an interface between the control system and another device and/or an external device interface. The control system may include at least one of a general purpose single- or multi-chip processor, a digital signal processor (DSP), an application specific integrated circuit (ASIC), a field programmable gate array (FPGA) or other programmable logic device, discrete gate or transistor logic, or discrete hardware components.

[0020] In some examples, the interface system may be capable of receiving audio data that may include a plurality of audio objects. The audio objects may include audio signals and associated audio object metadata. The audio object metadata may include at least audio object position metadata.

[0021] The control system may be capable of receiving reproduction environment data that may include an indication of a number of reproduction speakers in a reproduction environment. The control system may be capable of determining at least one audio object type from among a list of audio object types that may include dialogue. The control system may be capable of making an audio object prioritization based, at least in part, on the audio object type, and rendering the audio objects into a plurality of speaker feed signals based, at least in part, on the audio object position metadata. Making the audio object prioritization may involve assigning a highest priority to audio objects that correspond to dialogue.

However, in alternative implementations making the audio object prioritization may involve assigning a highest priority to audio objects that correspond to another audio object type. Each speaker feed signal may correspond to at least one of the reproduction speakers within the reproduction environment.

[0022] In some examples, the interface system may be capable of receiving hearing environment data. The hearing environment data may include at least one factor such as a model of hearing loss, a deficiency of at least one reproduction speaker and/or current environmental noise. The control system may be capable of adjusting the audio object levels based, at least in part, on the hearing environment data.

[0023] In some implementations, the control system may be capable of extracting one or more features from the audio data and determining an audio object type based, at least in part, on features extracted from the audio signals. The audio object type may be selected from a list of audio object types that includes dialogue. In some examples, the list of audio object types also may include background music, events and/or ambiance.

[0024] In some implementations, the control system may be capable of making an audio object prioritization based, at least in part, on the audio object type. The audio object

prioritization may determine, at least in part, a gain to be applied during a process of rendering the audio objects into speaker feed signals. In some examples, making the audio object prioritization may involve assigning a highest priority to audio objects that correspond to dialogue. However, in alternative implementations making the audio object prioritization may involve assigning a highest priority to audio objects that correspond to another audio object type. In some such examples, the control system may be capable of adding audio object prioritization metadata, based on the audio object prioritization, to the audio object metadata.

[0025] According to some examples, the interface system may be capable of receiving hearing environment data that may include a model of hearing loss. The hearing environment data may include environmental noise data, speaker deficiency data and/or hearing loss performance data. In some such examples, the control system may be capable of adjusting audio object levels according to the audio object prioritization and the hearing environment data. In some implementations, the control system may be capable of rendering the audio objects into a plurality of speaker feed signals based, at least in part, on the audio object position metadata. Each speaker feed signal may correspond to at least one of the reproduction speakers within the reproduction environment.

[0026] In some examples, the control system may include at least one excitation approximation module capable of determining excitation data. In some such examples, the excitation data may include an excitation indication (also referred to herein as an "excitation") for each of the plurality of audio objects. The excitation may be a function of a distribution of energy along a basilar membrane of a human ear. At least one of the excitations may be based, at least in part, on the hearing environment data. In some implementations, the control system may include a gain solver capable of receiving the excitation data and of determining gain data based, at least in part, on the excitations, the audio object prioritization and the hearing environment data.

[0027] Some or all of the methods described herein may be performed by one or more devices according to instructions (e.g., software) stored on non-transitory media. Such non-transitory media may include memory devices such as those described herein, including but not limited to random access memory (RAM) devices, read-only memory (ROM) devices, etc. Accordingly, some innovative aspects of the subject matter described in this disclosure can be implemented in a non-transitory medium having software stored thereon.

[0028] For example, the software may include instructions for controlling at least one device for receiving audio data that may include a plurality of audio objects. The audio

objects may include audio signals and associated audio object metadata. The audio object metadata may include at least audio object position metadata. In some examples, the software may include instructions for receiving reproduction environment data that may include an indication (direct and/or indirect) of a number of reproduction speakers in a reproduction environment.

5 [0029] In some such examples, the software may include instructions for determining at least one audio object type from among a list of audio object types that may include dialogue and for making an audio object prioritization based, at least in part, on the audio object type. In some implementations, making the audio object prioritization may involve  
10 assigning a highest priority to audio objects that correspond to dialogue. However, in alternative implementations making the audio object prioritization may involve assigning a highest priority to audio objects that correspond to another audio object type. In some such examples, the software may include instructions for adjusting audio object levels according to the audio object prioritization and for rendering the audio objects into a plurality of speaker  
15 feed signals based, at least in part, on the audio object position metadata. Each speaker feed signal may correspond to at least one of the reproduction speakers within the reproduction environment.

[0030] According to some examples, the software may include instructions for controlling the at least one device to receive hearing environment data that may include data  
20 corresponding to a model of hearing loss, a deficiency of at least one reproduction speaker and/or current environmental noise. Adjusting the audio object levels may be based, at least in part, on the hearing environment data.

[0031] In some implementations, the software may include instructions for extracting one or more features from the audio data and for determining an audio object type based, at  
25 least in part, on features extracted from the audio signals. The audio object type may be selected from a list of audio object types that includes dialogue.

[0032] In some such examples, the software may include instructions for making an audio object prioritization based, at least in part, on the audio object type. In some implementations, the audio object prioritization may determine, at least in part, a gain to be  
30 applied during a process of rendering the audio objects into speaker feed signals. Making the audio object prioritization may, in some examples, involve assigning a highest priority to audio objects that correspond to dialogue. However, in alternative implementations making the audio object prioritization may involve assigning a highest priority to audio objects that correspond to another audio object type. In some such examples, the software may include

instructions for adding audio object prioritization metadata, based on the audio object prioritization, to the audio object metadata.

[0033] According to some examples, the software may include instructions for controlling the at least one device to receive hearing environment data that may include data corresponding to a model of hearing loss, a deficiency of at least one reproduction speaker and/or current environmental noise. In some such examples, the software may include instructions for adjusting audio object levels according to the audio object prioritization and the hearing environment data and rendering the audio objects into a plurality of speaker feed signals based, at least in part, on the audio object position metadata. Each speaker feed signal may correspond to at least one of the reproduction speakers within the reproduction environment.

[0034] Details of one or more implementations of the subject matter described in this specification are set forth in the accompanying drawings and the description below. Other features, aspects, and advantages will become apparent from the description, the drawings, and the claims. Note that the relative dimensions of the following figures may not be drawn to scale.

### BRIEF DESCRIPTION OF THE DRAWINGS

[0035] Figure 1 shows an example of a playback environment having a Dolby Surround 5.1 configuration.

[0036] Figure 2 shows an example of a playback environment having a Dolby Surround 7.1 configuration.

[0037] Figures 3A and 3B illustrate two examples of home theater playback environments that include height speaker configurations.

[0038] Figure 4A shows an example of a graphical user interface (GUI) that portrays speaker zones at varying elevations in a virtual playback environment.

[0039] Figure 4B shows an example of another playback environment.

[0040] Figure 5A shows an example of an audio object and associated audio object width in a virtual reproduction environment.

[0041] Figure 5B shows an example of a spread profile corresponding to the audio object width shown in Figure 5A.

[0042] Figure 5C shows an example of virtual source locations relative to a playback environment.

[0043] Figure 5D shows an alternative example of virtual source locations relative to

a playback environment.

[0044] Figure 5E shows examples of W, X, Y and Z basis functions.

[0045] Figure 6A is a block diagram that represents some components that may be used for audio content creation.

5 [0046] Figure 6B is a block diagram that represents some components that may be used for audio playback in a reproduction environment (e.g., a movie theater).

[0047] Figure 7 is a block diagram that shows examples of components of an apparatus capable of implementing various aspects of this disclosure.

10 [0048] Figure 8 is a flow diagram that outlines one example of a method that may be performed by the apparatus of Figure 7.

[0049] Figure 9A is a block diagram that shows examples of an object prioritizer and an object renderer.

[0050] Figure 9B shows an example of object prioritizers and object Tenderers in two different contexts.

15 [0051] Figure 9C is a flow diagram that outlines one example of a method that may be performed by apparatus such as those shown in Figures, 7, 9A and/or 9B.

[0052] Figure 10 is a block diagram that shows examples of object prioritizer elements according to one implementation.

20 [0053] Figure 11 is a block diagram that shows examples of object renderer elements according to one implementation.

[0054] Figure 12 shows examples of dynamic range compression curves.

[0055] Figure 13 is a block diagram that illustrates examples of elements in a more detailed implementation.

25 [0056] Like reference numbers and designations in the various drawings indicate like elements.

## DESCRIPTION OF EXAMPLE EMBODIMENTS

[0057] The following description is directed to certain implementations for the purposes of describing some innovative aspects of this disclosure, as well as examples of contexts in which these innovative aspects may be implemented. However, the teachings herein can be applied in various different ways. For example, while various implementations are described in terms of particular playback environments, the teachings herein are widely applicable to other known playback environments, as well as playback environments that may be introduced in the future. Moreover, the described implementations may be

implemented, at least in part, in various devices and systems as hardware, software, firmware, cloud-based systems, etc. Accordingly, the teachings of this disclosure are not intended to be limited to the implementations shown in the figures and/or described herein, but instead have wide applicability.

5           **[0058]** As used herein, the term "audio object" refers to audio signals (also referred to herein as "audio object signals") and associated metadata that may be created or "authored" without reference to any particular playback environment. The associated metadata may include audio object position data, audio object gain data, audio object size data, audio object trajectory data, etc. As used herein, the term "rendering" refers to a process of transforming  
10 audio objects into speaker feed signals for a playback environment, which may be an actual playback environment or a virtual playback environment. A rendering process may be performed, at least in part, according to the associated metadata and according to playback environment data. The playback environment data may include an indication of a number of speakers in a playback environment and an indication of the location of each speaker within  
15 the playback environment.

**[0059]** Figure 1 shows an example of a playback environment having a Dolby Surround 5.1 configuration. In this example, the playback environment is a cinema playback environment. Dolby Surround 5.1 was developed in the 1990s, but this configuration is still widely deployed in home and cinema playback environments. In a cinema playback  
20 environment, a projector 105 may be configured to project video images, e.g. for a movie, on a screen 150. Audio data may be synchronized with the video images and processed by the sound processor 110. The power amplifiers 115 may provide speaker feed signals to speakers of the playback environment 100.

**[0060]** The Dolby Surround 5.1 configuration includes a left surround channel 120 for  
25 the left surround array 122 and a right surround channel 125 for the right surround array 127. The Dolby Surround 5.1 configuration also includes a left channel 130 for the left speaker array 132, a center channel 135 for the center speaker array 137 and a right channel 140 for the right speaker array 142. In a cinema environment, these channels may be referred to as a left screen channel, a center screen channel and a right screen channel, respectively. A  
30 separate low-frequency effects (LFE) channel 144 is provided for the subwoofer 145.

**[0061]** In 2010, Dolby provided enhancements to digital cinema sound by introducing Dolby Surround 7.1. Figure 2 shows an example of a playback environment having a Dolby Surround 7.1 configuration. A digital projector 205 may be configured to receive digital video data and to project video images on the screen 150. Audio data may be processed by

the sound processor 210. The power amplifiers 215 may provide speaker feed signals to speakers of the playback environment 200.

[0062] Like Dolby Surround 5.1, the Dolby Surround 7.1 configuration includes a left channel 130 for the left speaker array 132, a center channel 135 for the center speaker array 137, a right channel 140 for the right speaker array 142 and an LFE channel 144 for the subwoofer 145. The Dolby Surround 7.1 configuration includes a left side surround (Lss) array 220 and a right side surround (Rss) array 225, each of which may be driven by a single channel.

[0063] However, Dolby Surround 7.1 increases the number of surround channels by splitting the left and right surround channels of Dolby Surround 5.1 into four zones: in addition to the left side surround array 220 and the right side surround array 225, separate channels are included for the left rear surround (Lrs) speakers 224 and the right rear surround (Rrs) speakers 226. Increasing the number of surround zones within the playback environment 200 can significantly improve the localization of sound.

[0064] In an effort to create a more immersive environment, some playback environments may be configured with increased numbers of speakers, driven by increased numbers of channels. Moreover, some playback environments may include speakers deployed at various elevations, some of which may be "height speakers" configured to produce sound from an area above a seating area of the playback environment.

[0065] Figures 3A and 3B illustrate two examples of home theater playback environments that include height speaker configurations. In these examples, the playback environments 300a and 300b include the main features of a Dolby Surround 5.1 configuration, including a left surround speaker 322, a right surround speaker 327, a left speaker 332, a right speaker 342, a center speaker 337 and a subwoofer 145. However, the playback environment 300 includes an extension of the Dolby Surround 5.1 configuration for height speakers, which may be referred to as a Dolby Surround 5.1.2 configuration.

[0066] Figure 3A illustrates an example of a playback environment having height speakers mounted on a ceiling 360 of a home theater playback environment. In this example, the playback environment 300a includes a height speaker 352 that is in a left top middle (Ltm) position and a height speaker 357 that is in a right top middle (Rtm) position. In the example shown in Figure 3B, the left speaker 332 and the right speaker 342 are Dolby Elevation speakers that are configured to reflect sound from the ceiling 360. If properly configured, the reflected sound may be perceived by listeners 365 as if the sound source originated from the ceiling 360. However, the number and configuration of speakers is

merely provided by way of example. Some current home theater implementations provide for up to 34 speaker positions, and contemplated home theater implementations may allow yet more speaker positions.

5 [0067] Accordingly, the modern trend is to include not only more speakers and more channels, but also to include speakers at differing heights. As the number of channels increases and the speaker layout transitions from Y2D to 3D, the tasks of positioning and rendering sounds becomes increasingly difficult.

10 [0068] Accordingly, Dolby has developed various tools, including but not limited to user interfaces, which increase functionality and/or reduce authoring complexity for a 3D audio sound system. Some such tools may be used to create audio objects and/or metadata for audio objects.

15 [0069] Figure 4A shows an example of a graphical user interface (GUI) that portrays speaker zones at varying elevations in a virtual playback environment. GUI 400 may, for example, be displayed on a display device according to instructions from a logic system, according to signals received from user input devices, etc. Some such devices are described below with reference to Figure 11.

20 [0070] As used herein with reference to virtual playback environments such as the virtual playback environment 404, the term "speaker zone" generally refers to a logical construct that may or may not have a one-to-one correspondence with a speaker of an actual playback environment. For example, a "speaker zone location" may or may not correspond to a particular speaker location of a cinema playback environment. Instead, the term "speaker zone location" may refer generally to a zone of a virtual playback environment. In some implementations, a speaker zone of a virtual playback environment may correspond to a virtual speaker, e.g., via the use of virtualizing technology such as Dolby Headphone,<sup>TM</sup>  
25 (sometimes referred to as Mobile Surround<sup>TM</sup>), which creates a virtual surround sound environment in real time using a set of two-channel stereo headphones. In GUI 400, there are seven speaker zones 402a at a first elevation and two speaker zones 402b at a second elevation, making a total of nine speaker zones in the virtual playback environment 404. In this example, speaker zones 1-3 are in the front area 405 of the virtual playback environment  
30 404. The front area 405 may correspond, for example, to an area of a cinema playback environment in which a screen 150 is located, to an area of a home in which a television screen is located, etc.

[0071] Here, speaker zone 4 corresponds generally to speakers in the left area 410 and speaker zone 5 corresponds to speakers in the right area 415 of the virtual playback

environment 404. Speaker zone 6 corresponds to a left rear area 412 and speaker zone 7 corresponds to a right rear area 414 of the virtual playback environment 404. Speaker zone 8 corresponds to speakers in an upper area 420a and speaker zone 9 corresponds to speakers in an upper area 420b, which may be a virtual ceiling area. Accordingly, the locations of speaker zones 1-9 that are shown in Figure 4A may or may not correspond to the locations of speakers of an actual playback environment. Moreover, other implementations may include more or fewer speaker zones and/or elevations.

[0072] In various implementations described herein, a user interface such as GUI 400 may be used as part of an authoring tool and/or a rendering tool. In some implementations, the authoring tool and/or rendering tool may be implemented via software stored on one or more non-transitory media. The authoring tool and/or rendering tool may be implemented (at least in part) by hardware, firmware, etc., such as the logic system and other devices described below with reference to Figure 11. In some authoring implementations, an associated authoring tool may be used to create metadata for associated audio data. The metadata may, for example, include data indicating the position and/or trajectory of an audio object in a three-dimensional space, speaker zone constraint data, etc. The metadata may be created with respect to the speaker zones 402 of the virtual playback environment 404, rather than with respect to a particular speaker layout of an actual playback environment. A rendering tool may receive audio data and associated metadata, and may compute audio gains and speaker feed signals for a playback environment. Such audio gains and speaker feed signals may be computed according to an amplitude panning process, which can create a perception that a sound is coming from a position  $P$  in the playback environment. For example, speaker feed signals may be provided to speakers 1 through  $N$  of the playback environment according to the following equation:

$$[0073] \quad x_i(t) = g_i x(t), \quad i = 1, \dots, N \quad (\text{Equation 1})$$

[0074] In Equation 1,  $x_i(t)$  represents the speaker feed signal to be applied to speaker  $i$ ,  $g_i$  represents the gain factor of the corresponding channel,  $x(t)$  represents the audio signal and  $t$  represents time. The gain factors may be determined, for example, according to the amplitude panning methods described in Section 2, pages 3-4 of V. Pulkki, *Compensating Displacement of Amplitude-Panned Virtual Sources* (Audio Engineering Society (AES) International Conference on Virtual, Synthetic and Entertainment Audio), which is hereby incorporated by reference. In some implementations, the gains may be frequency dependent. In some implementations, a time delay may be introduced by replacing  $x(t)$  by  $x(t-At)$ .

[0075] In some rendering implementations, audio reproduction data created with

reference to the speaker zones 402 may be mapped to speaker locations of a wide range of playback environments, which may be in a Dolby Surround 5.1 configuration, a Dolby Surround 7.1 configuration, a Hamasaki 22.2 configuration, or another configuration. For example, referring to Figure 2, a rendering tool may map audio reproduction data for speaker zones 4 and 5 to the left side surround array 220 and the right side surround array 225 of a playback environment having a Dolby Surround 7.1 configuration. Audio reproduction data for speaker zones 1, 2 and 3 may be mapped to the left screen channel 230, the right screen channel 240 and the center screen channel 235, respectively. Audio reproduction data for speaker zones 6 and 7 may be mapped to the left rear surround speakers 224 and the right rear surround speakers 226.

[0076] Figure 4B shows an example of another playback environment. In some implementations, a rendering tool may map audio reproduction data for speaker zones 1, 2 and 3 to corresponding screen speakers 455 of the playback environment 450. A rendering tool may map audio reproduction data for speaker zones 4 and 5 to the left side surround array 460 and the right side surround array 465 and may map audio reproduction data for speaker zones 8 and 9 to left overhead speakers 470a and right overhead speakers 470b. Audio reproduction data for speaker zones 6 and 7 may be mapped to left rear surround speakers 480a and right rear surround speakers 480b.

[0077] In some authoring implementations, an authoring tool may be used to create metadata for audio objects. The metadata may indicate the 3D position of the object, rendering constraints, content type (e.g. dialog, effects, etc.) and/or other information. Depending on the implementation, the metadata may include other types of data, such as width data, gain data, trajectory data, etc. Some audio objects may be static, whereas others may move.

[0078] Audio objects are rendered according to their associated metadata, which generally includes positional metadata indicating the position of the audio object in a three-dimensional space at a given point in time. When audio objects are monitored or played back in a playback environment, the audio objects are rendered according to the positional metadata using the speakers that are present in the playback environment, rather than being output to a predetermined physical channel, as is the case with traditional, channel-based systems such as Dolby 5.1 and Dolby 7.1.

[0079] In addition to positional metadata, other types of metadata may be necessary to produce intended audio effects. For example, in some implementations, the metadata associated with an audio object may indicate audio object size, which may also be referred to

as "width." Size metadata may be used to indicate a spatial area or volume occupied by an audio object. A spatially large audio object should be perceived as covering a large spatial area, not merely as a point sound source having a location defined only by the audio object position metadata. In some instances, for example, a large audio object should be perceived as occupying a significant portion of a playback environment, possibly even surrounding the listener.

[0080] Spread and apparent source width control are features of some existing surround sound authoring/rendering systems. In this disclosure, the term "spread" refers to distributing the same signal over multiple speakers to blur the sound image. The term "width" (also referred to herein as "size" or "audio object size") refers to decorrelating the output signals to each channel for apparent width control. Width may be an additional scalar value that controls the amount of decorrelation applied to each speaker feed signal.

[0081] Some implementations described herein provide a 3D axis oriented spread control. One such implementation will now be described with reference to Figures 5A and 5B. Figure 5A shows an example of an audio object and associated audio object width in a virtual reproduction environment. Here, the GUI 400 indicates an ellipsoid 555 extending around the audio object 510, indicating the audio object width or size. The audio object width may be indicated by audio object metadata and/or received according to user input. In this example, the x and y dimensions of the ellipsoid 555 are different, but in other implementations these dimensions may be the same. The z dimensions of the ellipsoid 555 are not shown in Figure 5A.

[0082] Figure 5B shows an example of a spread profile corresponding to the audio object width shown in Figure 5A. Spread may be represented as a three-dimensional vector parameter. In this example, the spread profile 507 can be independently controlled along 3 dimensions, e.g., according to user input. The gains along the x and y axes are represented in Figure 5B by the respective height of the curves 560 and 1520. The gain for each sample 562 is also indicated by the size of the corresponding circles 575 within the spread profile 507. The responses of the speakers 580 are indicated by gray shading in Figure 5B.

[0083] In some implementations, the spread profile 507 may be implemented by a separable integral for each axis. According to some implementations, a minimum spread value may be set automatically as a function of speaker placement to avoid timbral discrepancies when panning. Alternatively, or additionally, a minimum spread value may be set automatically as a function of the velocity of the panned audio object, such that as audio object velocity increases an object becomes more spread out spatially, similarly to how

rapidly moving images in a motion picture appear to blur.

[0084] The human hearing system is very sensitive to changes in the correlation or coherence of the signals arriving at both ears, and maps this correlation to a perceived object size attribute if the normalized correlation is smaller than the value of +1. Therefore, in order to create a convincing spatial object size, or spatial diffuseness, a significant proportion of the speaker signals in a playback environment should be mutually independent, or at least be uncorrected (e.g. independent in terms of first-order cross correlation or covariance). A satisfactory decorrelation process is typically rather complex, normally involving time-variant filters.

[0085] A cinema sound track may include hundreds of objects, each with its associated position metadata, size metadata and possibly other spatial metadata. Moreover, a cinema sound system can include hundreds of loudspeakers, which may be individually controlled to provide satisfactory perception of audio object locations and sizes. In a cinema, therefore, hundreds of objects may be reproduced by hundreds of loudspeakers, and the object-to-loudspeaker signal mapping consists of a very large matrix of panning coefficients. When the number of objects is given by  $M$ , and the number of loudspeakers is given by  $N$ , this matrix has up to  $M*N$  elements. This has implications for the reproduction of diffuse or large-size objects. In order to create a convincing spatial object size, or spatial diffuseness, a significant proportion of the  $N$  loudspeaker signals should be mutually independent, or at least be uncorrected. This generally involves the use of many (up to  $N$ ) independent decorrelation processes, causing a significant processing load for the rendering process. Moreover, the amount of decorrelation may be different for each object, which further complicates the rendering process. A sufficiently complex rendering system, such as a rendering system for a commercial theater, may be capable of providing such decorrelation.

[0086] However, less complex rendering systems, such as those intended for home theater systems, may not be capable of providing adequate decorrelation. Some such rendering systems are not capable of providing decorrelation at all. Decorrelation programs that are simple enough to be executed on a home theater system can introduce artifacts. For example, comb-filter artifacts may be introduced if a low-complexity decorrelation process is followed by a downmix process.

[0087] Another potential problem is that in some applications, object-based audio is transmitted in the form of a backward-compatible mix (such as Dolby Digital or Dolby Digital Plus), augmented with additional information for retrieving one or more objects from that backward-compatible mix. The backward-compatible mix would normally not have the

effect of decorrelation included. In some such systems, the reconstruction of objects may only work reliably if the backward-compatible mix was created using simple panning procedures. The use of decorrelators in such processes can harm the audio object reconstruction process, sometimes severely. In the past, this has meant that one could either  
5 choose not to apply decorrelation in the backward-compatible mix, thereby degrading the artistic intent of that mix, or accept degradation in the object reconstruction process.

[0088] In order to address such potential problems, some implementations described herein involve identifying diffuse or spatially large audio objects for special processing. Such methods and devices may be particularly suitable for audio data to be rendered in a  
10 home theater. However, these methods and devices are not limited to home theater use, but instead have broad applicability.

[0089] Due to their spatially diffuse nature, objects with a large size are not perceived as point sources with a compact and concise location. Therefore, multiple speakers are used to reproduce such spatially diffuse objects. However, the exact locations of the speakers in  
15 the playback environment that are used to reproduce large audio objects are less critical than the locations of speakers use to reproduce compact, small-sized audio objects. Accordingly, a high-quality reproduction of large audio objects is possible without prior knowledge about the actual playback speaker configuration used to eventually render decorrelated large audio object signals to actual speakers of the playback environment. Consequently, decorrelation  
20 processes for large audio objects can be performed "upstream," before the process of rendering audio data for reproduction in a playback environment, such as a home theater system, for listeners. In some examples, decorrelation processes for large audio objects are performed prior to encoding audio data for transmission to such playback environments.

[0090] Such implementations do not require the renderer of a playback environment  
25 to be capable of high-complexity decorrelation, thereby allowing for rendering processes that may be relatively simpler, more efficient and cheaper. Backward-compatible downmixes may include the effect of decorrelation to maintain the best possible artistic intent, without the need to reconstruct the object for rendering-side decorrelation. High-quality decorrelators can be applied to large audio objects upstream of a final rendering process, e.g., during an  
30 authoring or post-production process in a sound studio. Such decorrelators may be robust with regard to downmixing and/or other downstream audio processing.

[0091] Some examples of rendering audio object signals to virtual speaker locations will now be described with reference to Figures 5C and 5D. Figure 5C shows an example of virtual source locations relative to a playback environment. The playback environment may

be an actual playback environment or a virtual playback environment. The virtual source locations 505 and the speaker locations 525 are merely examples. However, in this example the playback environment is a virtual playback environment and the speaker locations 525 correspond to virtual speaker locations.

5           **[0092]** In some implementations, the virtual source locations 505 may be spaced uniformly in all directions. In the example shown in Figure 5A, the virtual source locations 505 are spaced uniformly along x, y and z axes. The virtual source locations 505 may form a rectangular grid of  $N_x$  by  $N_y$  by  $N_z$  virtual source locations 505. In some implementations, the value of  $N$  may be in the range of 5 to 100. The value of  $N$  may depend, at least in part, on  
10 the number of speakers in the playback environment (or expected to be in the playback environment): it may be desirable to include two or more virtual source locations 505 between each speaker location.

**[0093]** However, in alternative implementations, the virtual source locations 505 may be spaced differently. For example, in some implementations the virtual source locations 505  
15 may have a first uniform spacing along the x and y axes and a second uniform spacing along the z axis. In other implementations, the virtual source locations 505 may be spaced non-uniformly.

**[0094]** In this example, the audio object volume 520a corresponds to the size of the audio object. The audio object 510 may be rendered according to the virtual source locations  
20 505 enclosed by the audio object volume 520a. In the example shown in Figure 5A, the audio object volume 520a occupies part, but not all, of the playback environment 500a. Larger audio objects may occupy more of (or all of) the playback environment 500a. In some examples, if the audio object 510 corresponds to a point source, the audio object 510 may have a size of zero and the audio object volume 520a may be set to zero.

25           **[0095]** According to some such implementations, an authoring tool may link audio object size with decorrelation by indicating (e.g., via a decorrelation flag included in associated metadata) that decorrelation should be turned on when the audio object size is greater than or equal to a size threshold value and that decorrelation should be turned off if the audio object size is below the size threshold value. In some implementations,  
30 decorrelation may be controlled (e.g., increased, decreased or disabled) according to user input regarding the size threshold value and/or other input values.

**[0096]** In this example, the virtual source locations 505 are defined within a virtual source volume 502. In some implementations, the virtual source volume may correspond with a volume within which audio objects can move. In the example shown in Figure 5A, the

playback environment 500a and the virtual source volume 502a are co-extensive, such that each of the virtual source locations 505 corresponds to a location within the playback environment 500a. However, in alternative implementations, the playback environment 500a and the virtual source volume 502 may not be co-extensive.

5           **[0097]** For example, at least some of the virtual source locations 505 may correspond to locations outside of the playback environment. Figure 5B shows an alternative example of virtual source locations relative to a playback environment. In this example, the virtual source volume 502b extends outside of the playback environment 500b. Some of the virtual source locations 505 within the audio object volume 520b are located inside of the playback environment 500b and other virtual source locations 505 within the audio object volume 520b are located outside of the playback environment 500b.

**[0098]** In other implementations, the virtual source locations 505 may have a first uniform spacing along x and y axes and a second uniform spacing along a z axis. The virtual source locations 505 may form a rectangular grid of  $N_x$  by  $N_y$  by  $M_z$  virtual source locations  
15 505. For example, in some implementations there may be fewer virtual source locations 505 along the z axis than along the x or y axes. In some such implementations, the value of  $N$  may be in the range of 10 to 100, whereas the value of  $M$  may be in the range of 5 to 10.

**[0099]** Some implementations involve computing gain values for each of the virtual source locations 505 within an audio object volume 520. In some implementations, gain  
20 values for each channel of a plurality of output channels of a playback environment (which may be an actual playback environment or a virtual playback environment) will be computed for each of the virtual source locations 505 within an audio object volume 520. In some implementations, the gain values may be computed by applying a vector-based amplitude panning ("VBAP") algorithm, a pairwise panning algorithm or a similar algorithm to  
25 compute gain values for point sources located at each of the virtual source locations 505 within an audio object volume 520. In other implementations, a separable algorithm, to compute gain values for point sources located at each of the virtual source locations 505 within an audio object volume 520. As used herein, a "separable" algorithm is one for which the gain of a given speaker can be expressed as a product of multiple factors (e.g., three  
30 factors), each of which depends only on one of the coordinates of the virtual source location 505. Examples include algorithms implemented in various existing mixing console panners, including but not limited to the Pro Tools™ software and panners implemented in digital film consoles provided by AMS Neve.

**[00100]** In some implementations, a virtual acoustic space may be represented

as an approximation to the sound field at a point (or on a sphere). Some such implementations may involve projecting onto a set of orthogonal basis functions on a sphere. In some such representations, which are based on Ambisonics, the basis functions are spherical harmonics. In such a format, a source at azimuth angle  $\Theta$  and an elevation  $\phi$  will be panned with different gains onto the first 4 W, X, Y and Z basis functions. In some such examples, the gains may be given by the following equations:

$$W = S \cdot \frac{1}{\sqrt{2}}$$

$$X = S \cdot \cos \Theta \cos \phi$$

$$Y = S \cdot \sin \Theta \cos \phi$$

$$Z = S \cdot \sin \phi$$

10

[00101] Figure 5E shows examples of W, X, Y and Z basis functions. In this example, the omnidirectional component W is independent of angle. The X, Y and Z components may, for example, correspond to microphones with a dipole response, oriented along the X, Y and Z axes. Higher order components, examples of which are shown in rows 15 550 and 555 of Figure 5E, can be used to achieve greater spatial accuracy.

[00102] Mathematically, the spherical harmonics are solutions of Laplace's equation in 3 dimensions, and are found to have the form  $Y_l^m(\theta, \phi) = N Y_l^m(\cos \theta) P_l^m(\cos \phi)$ , in which  $m$  represents an integer,  $N$  represents a normalization constant and  $P_l^m$  represents a Legendre polynomial. However, in some implementations the above functions may be represented in rectangular coordinates rather than the spherical coordinates used above.

20

[00103] Figure 6A is a block diagram that represents some components that may be used for audio content creation. The system 600 may, for example, be used for audio content creation in mixing studios and/or dubbing stages. In this example, the system 600 includes an audio and metadata authoring tool 605 and a rendering tool 610. In this implementation, the audio and metadata authoring tool 605 and the rendering tool 610 include audio connect interfaces 607 and 612, respectively, which may be configured for communication via AES/EBU, MADI, analog, etc. The audio and metadata authoring tool 605 and the rendering tool 610 include network interfaces 609 and 617, respectively, which may be configured to send and receive metadata via TCP/IP or any other suitable protocol. The interface 620 is configured to output audio data to speakers.

30

[00104] The system 600 may, for example, include an existing authoring system, such as a Pro Tools™ system, running a metadata creation tool (i.e., a panner as

described herein) as a plugin. The panner could also run on a standalone system (e.g. a PC or a mixing console) connected to the rendering tool 610, or could run on the same physical device as the rendering tool 610. In the latter case, the panner and renderer could use a local connection e.g., through shared memory. The panner GUI could also be remoted on a tablet  
5 device, a laptop, etc. The rendering tool 610 may comprise a rendering system that includes a sound processor capable of executing rendering software. The rendering system may include, for example, a personal computer, a laptop, etc., that includes interfaces for audio input/output and an appropriate logic system.

[00105] Figure 6B is a block diagram that represents some components that  
10 may be used for audio playback in a reproduction environment (e.g., a movie theater). The system 650 includes a cinema server 655 and a rendering system 660 in this example. The cinema server 655 and the rendering system 660 include network interfaces 657 and 662, respectively, which may be configured to send and receive audio objects via TCP/IP or any other suitable protocol. The interface 664 may be configured to output audio data to  
15 speakers.

[00106] As noted above, it can be challenging for listeners with hearing loss to  
hear all sounds that are reproduced during a movie, a television program, etc. For example, listeners with hearing loss may perceive an audio scene (the aggregate of audio objects being reproduced at a particular time) as seeming to be too "cluttered," in other words having too  
20 many audio objects. It may be difficult for listeners with hearing loss to understand dialogue, for example. Listeners who have normal hearing can experience similar difficulties in a noisy playback environment.

[00107] Some implementations disclosed herein provide methods for  
improving an audio scene for people suffering from hearing loss or for adverse hearing  
25 environments. Some such implementations are based, at least in part, on the observation that some audio objects may be more important to an audio scene than others. Accordingly, in some such implementations audio objects may be prioritized. For example, in some implementations, audio objects that correspond to dialogue may be assigned the highest priority. Other implementations may involve assigning the highest priority to other types of  
30 audio objects, such as audio objects that correspond to events. In some examples, during a process of dynamic range compression, higher-priority audio objects may be boosted more, or cut less, than lower-priority audio objects. Lower-priority audio objects may fall completely below the threshold of hearing, in which case they may be dropped and not rendered.

[00108] Figure 7 is a block diagram that shows examples of components of an apparatus capable of implementing various aspects of this disclosure. The apparatus 700 may be implemented via hardware, via software stored on non-transitory media, via firmware and/or by combinations thereof. The types and numbers of components shown in Figure 7 are merely shown by way of example. Alternative implementations may include more, fewer and/or different components. The apparatus 700 may, for example, be an instance of an apparatus such as those described below with reference to Figures 8-13. In some examples, the apparatus 700 may be a component of another device or of another system. For example, the apparatus 700 may be a component of an authoring system such as the system 600 described above or a component of a system used for audio playback in a reproduction environment (e.g., a movie theater, a home theater system, etc.) such as the system 650 described above.

[00109] In this example, the apparatus 700 includes an interface system 705 and a control system 710. The interface system 705 may include one or more network interfaces, one or more interfaces between the control system 710 and a memory system and/or one or more an external device interfaces (such as one or more universal serial bus (USB) interfaces). The control system 710 may, for example, include a general purpose single- or multi-chip processor, a digital signal processor (DSP), an application specific integrated circuit (ASIC), a field programmable gate array (FPGA) or other programmable logic device, discrete gate or transistor logic, and/or discrete hardware components. In some implementations, the control system 710 may be capable of authoring system functionality and/or audio playback functionality.

[00110] Figure 8 is a flow diagram that outlines one example of a method that may be performed by the apparatus of Figure 7. The blocks of method 800, like other methods described herein, are not necessarily performed in the order indicated. Moreover, such methods may include more or fewer blocks than shown and/or described.

[00111] In this implementation, block 805 involves receiving audio data that includes a plurality of audio objects. In this example, the audio objects include audio signals (which may also be referred to herein as "audio object signals") and associated audio object metadata. In this implementation, the audio object metadata includes audio object position metadata. In some implementations, the audio object metadata may include one or more other types of audio object metadata, such as audio object type metadata, audio object size metadata, audio object prioritization metadata and/or one or more other types of audio object metadata.

[00112] In the example shown in Figure 8, block 810 involves receiving reproduction environment data. Here, the reproduction environment data includes an indication of a number of reproduction speakers in a reproduction environment. In some examples, positions of reproduction speakers in the reproduction environment may be determined, or inferred, according to the reproduction environment configuration. Accordingly, the reproduction environment data may or may not include an express indication of positions of reproduction speakers in the reproduction environment. In some implementations, the reproduction environment may be an actual reproduction environment, whereas in other implementations the reproduction environment may be a virtual reproduction environment.

[00113] In some examples, the reproduction environment data may include an indication of positions of reproduction speakers in the reproduction environment. In some implementations, the reproduction environment data may include an indication of a reproduction environment configuration. For example, the reproduction environment data may indicate whether the reproduction environment has a Dolby Surround 5.1 configuration, a Dolby Surround 7.1 configuration, a Hamasaki 22.2 surround sound configuration, a headphone configuration, a Dolby Surround 5.1.2 configuration, a Dolby Surround 7.1.2 configuration, a Dolby Atmos configuration or another reproduction environment configuration.

[00114] In this implementation, block 815 involves determining at least one audio object type from among a list of audio object types that includes dialogue. For example, a dialogue audio object may correspond to the speech of a particular individual. In some examples, the list of audio object types may include background music, events and/or ambiance. As noted above, in some instances the audio object metadata may include audio object type metadata. According to some such implementations, determining the audio object type may involve evaluating the object type metadata. Alternatively, or additionally, determining the audio object type may involve analyzing the audio signals of audio objects, e.g., as described below.

[00115] In the example shown in Figure 8, block 820 involves making an audio object prioritization based, at least in part, on the audio object type. In this implementation, making the audio object prioritization involves assigning a highest priority to audio objects that correspond to dialogue. In alternative implementations, making the audio object prioritization may involve assigning a highest priority to audio objects according to one or more other attributes, such as audio object volume or level. According

to some implementations, some events (such as explosions, bullets sounds, etc.) may be assigned a higher priority than dialogue and other events (such as the sounds of a fire) may be assigned a lower priority than dialogue.

5           **[00116]**           In some examples, the audio object metadata may include audio object size metadata. In some implementations, making an audio object prioritization may involve assigning a relatively lower priority to large or diffuse audio objects. For example, making the audio object prioritization may involve applying a function that reduces a priority of at least some audio objects (e.g., of non-dialogue audio objects) according to increases in audio object size. In some implementations, the function may not reduce the priority of audio  
10 objects that are below a threshold size.

**[00117]**           In this implementation, block 825 involves adjusting audio object levels according to the audio object prioritization. If the audio object metadata includes audio object prioritization metadata, adjusting the audio object levels may be based, at least in part, on the audio object prioritization metadata. In some implementations, the process of  
15 adjusting the audio object levels may be performed on multiple frequency bands of audio signals corresponding to an audio object. Adjusting the audio object levels may involves differentially adjusting levels of various frequency bands. However, in some implementations the process of adjusting the audio object levels may involve determining a single level adjustment for multiple frequency bands.

20           **[00118]**           Some instances may involve selecting at least one audio object that will not be rendered based, at least in part, on the audio object prioritization. According to some such examples, adjusting the audio object's level(s) according to the audio object prioritization may involve adjusting the audio object's level(s) such that the audio object's level(s) fall completely below the normal thresholds of human hearing, or below a particular  
25 listener's threshold of hearing. In some such examples, the audio object may be discarded and not rendered.

**[00119]**           According to some implementations, adjusting the audio object levels may involve dynamic range compression and/or automatic gain control processes. In some examples, during a process of dynamic range compression, the levels of higher-priority  
30 objects may be boosted more, or cut less, than the levels of lower-priority objects.

**[00120]**           Some implementations may involve receiving hearing environment data. In some such implementations, the hearing environment data may include a model of hearing loss, data corresponding to a deficiency of at least one reproduction speaker and/or data corresponding to current environmental noise. According to some such

implementations, adjusting the audio object levels may be based, at least in part, on the hearing environment data.

[00121] In the example shown in Figure 8, block 830 involves rendering the audio objects into a plurality of speaker feed signals based, at least in part, on the audio object position metadata, wherein each speaker feed signal corresponds to at least one of the reproduction speakers within the reproduction environment. In some examples, the reproduction speakers may be headphone speakers. The reproduction environment may be an actual acoustic space or a virtual acoustic space, depending on the particular implementation.

[00122] Accordingly, in some implementations block 830 may involve rendering the audio objects to locations in a virtual acoustic space. In some examples, block 830 may involve rendering the audio objects according to a plurality of virtual speaker locations within a virtual acoustic space. As described in more detail below, some examples may involve increasing a distance between at least some audio objects in the virtual acoustic space. In some instances, the virtual acoustic space may include a front area and a back area. The front area and the back area may, for example, be determined relative to a position of a virtual listener's head in the virtual acoustic space. The rendering may involve increasing a distance between at least some audio objects in the front area of the virtual acoustic space. Increasing this distance may, in some examples, improve the ability of a listener to hear the rendered audio objects more clearly. For example, increasing this distance may make dialogue more intelligible for some listeners.

[00123] In some implementations wherein a virtual acoustic space is represented by spherical harmonics (such as the implementations described above with reference to Figure 5B), the angular separation (as indicated by angle  $\theta$  and/or  $\phi$ ) between at least some audio objects in the front area of the virtual acoustic space may be increased prior to a rendering process. In some such implementations, the azimuthal angle  $\theta$  may be "warped" in such a way that at least some angles corresponding to an area in front of the virtual listener's head may be increased and at least some angles corresponding to an area behind the virtual listener's head may be decreased.

[00124] Some implementations may involve determining whether an audio object has audio signals that include a directional component and a diffuse component. If it is determined that the audio object has audio signals that include a directional component and a diffuse component, such implementations may involve reducing a level of the diffuse component.

[00125] For example, referring to Figures 5A and 5B, a single audio object

510 may include a plurality of gains, each of which may correspond with a different position in an actual or virtual space that is within an area or volume of the audio object 510. In Figure 5B, the gain for each sample 562 is indicated by the size of the corresponding circles 575 within the spread profile 507. The responses of the speakers 580 (which may be real speakers or virtual speakers) are indicated by gray shading in Figure 5B. In some implementations, gains corresponding to a position at or near the center of the ellipsoid 555 (e.g., the gain represented by the circle 575a) may correspond with a directional component of the audio signals, whereas gains corresponding to other positions within the ellipsoid 555 (e.g., the gain represented by the circle 575b) may correspond with a diffuse component of the audio signals.

[00126] Similarly, referring to Figures 5C and 5D, the audio object volumes 520a and 520b correspond to the size of the corresponding audio object 510. In some implementations, the audio object 510 may be rendered according to the virtual source locations 505 enclosed by the audio object volume 520a or 520b. In some such implementations, the audio object 510 may have a directional component associated with the position 515, which is in the center of the audio object volumes in these examples, and may have diffuse components associated with other virtual source locations 505 enclosed by the audio object volume 520a or 520b.

[00127] However, in some implementations an audio object's audio signals may include diffuse components that may not directly correspond to audio object size. For example, some such diffuse components may correspond to simulated reverberation, wherein the sound of an audio object source is reflected from various surfaces (such as walls) of a simulated room. By reducing these diffuse components of the audio signals, one can reduce the amount of room reverberation and produce a less "cluttered" sound.

[00128] Figure 9A is a block diagram that shows examples of an object prioritizer and an object renderer. The apparatus 900 may be implemented via hardware, via software stored on non-transitory media, via firmware and/or by combinations thereof. In some implementations, the apparatus 900 may be implemented in an authoring/content creation context, such as an audio editing context for a video, for a movie, for a game, etc. However, in other implementations the apparatus 900 may be implemented in a cinema context, a home theater context, or another consumer-related context.

[00129] In this example, the object prioritizer 905 is capable of making an audio object prioritization based, at least in part, on audio object type. For example, in some implementations, the object prioritizer 905 may assign the highest priority to audio objects

that correspond to dialogue. In other implementations, the object prioritizer 905 may assign the highest priority to other types of audio objects, such as audio objects that correspond to events. In some examples, more than one audio object may be assigned the same level of priority. For instance, two audio objects that correspond to dialogue may both be assigned the same priority. In this example, the object prioritizer 905 is capable of providing audio object prioritization metadata to the object renderer 910.

**[00130]** In some examples, the audio object type may be indicated by audio object metadata received by the object prioritizer 905. Alternatively, or additionally, in some implementations the object prioritizer 905 may be capable of making an audio object type determination based, at least in part, on an analysis of audio signals corresponding to audio objects. For example, according to some such implementations the object prioritizer 905 may be capable of making an audio object type determination based, at least in part, on features extracted from the audio signals. In some such implementations, the object prioritizer 905 may include a feature detector and a classifier. One such example is described below with reference to Figure 10.

**[00131]** According to some examples, the object prioritizer 905 may determine priority based, at least in part, on loudness and/or audio object size. For example, the object prioritizer 905 may indicate a relatively higher priority to relatively louder audio objects. In some instances, the object prioritizer 905 may assign a relatively lower priority to relatively larger audio objects. In some such examples, large audio objects (e.g., audio object having a size that is greater than a threshold size) may be assigned a relatively low priority unless the audio object is loud (e.g., has a loudness that is greater than a threshold level). Additional examples of object prioritization functionality are disclosed herein, including but not limited to those provided by Figure 10 and the corresponding description.

**[00132]** In the example shown in Figure 9A, the object prioritizer 905 may be capable of receiving user input. Such user input may, for example, be received via a user input system of the apparatus 900. The user input system may include a touch sensor or gesture sensor system and one or more associated controllers, a microphone for receiving voice commands and one or more associated controllers, a display and one or more associated controllers for providing a graphical user interface, etc. The controllers may be part of a control system such as the control system 710 that is shown in Figure 7 and described above. However, in some examples one or more of the controllers may reside in another device. For example, one or more of the controllers may reside in a server that is capable of providing voice activity detection functionality.

[00133] In some such implementations, a prioritization method applied by the object prioritizer 905 may be based, at least in part, on such user input. For example, in some implementations the type of audio object that will be assigned the highest priority may be determined according to user input. According to some examples, the priority level of selected audio objects may be determined according to user input. Such capabilities may, for example, be useful in the content creation/authoring context, e.g., for post-production editing of the audio for a movie, a video, etc. In some implementations, the number of priority levels in a hierarchy of priorities may be changed according to user input. For example, some such implementations may have a "default" number of priority levels (such as three levels corresponding to a highest level, a middle level and a lowest level). In some implementations, the number of priority levels may be increased or decreased according to user input (e.g., from 3 levels to 5 levels, from 4 levels to 3 levels, etc.).

[00134] In the implementation shown in Figure 9A, the object renderer 910 is capable of generating speaker feed signals for a reproduction environment based on received hearing environment data, audio signals and audio object metadata. The reproduction environment may be a virtual reproduction environment or an actual reproduction environment, depending on the particular implementation. In this example, the audio object metadata includes audio object prioritization metadata that is received from the object prioritizer 905. As noted elsewhere herein, the renderer may generate the speaker feed signals according to a particular reproduction environment configuration, which may be a headphone configuration, a non-headphone stereo configuration, a Dolby Surround 5.1 configuration, a Dolby Surround 7.1 configuration, a Hamasaki 22.2 surround sound configuration, a Dolby Surround 5.1.2 configuration, a Dolby Surround 7.1.2 configuration, a Dolby Atmos configuration, or some other configuration.

[00135] In some implementations the object renderer 910 may be capable of rendering audio objects to locations in a virtual acoustic space. In some such examples the object renderer 910 may be capable of increasing a distance between at least some audio objects in the virtual acoustic space. In some instances, the virtual acoustic space may include a front area and a back area. The front area and the back area may, for example, be determined relative to a position of a virtual listener's head in the virtual acoustic space. In some implementations, the object renderer 910 may be capable of increasing a distance between at least some audio objects in the front area of the virtual acoustic space.

[00136] The hearing environment data may include a model of hearing loss. According to some implementations, such a model may be an audiogram of a particular

individual, based on a hearing examination. Alternatively, or additionally, the hearing loss model may be a statistical model based on empirical hearing loss data for many individuals. In some examples, hearing environment data may include a function that may be used to calculate loudness (e.g., per frequency band) based on excitation level.

5           **[00137]**       In some instances, the hearing environment data may include data regarding a characteristic (e.g., a deficiency) of at least one reproduction speaker. Some speakers may, for example, distort when driven at particular frequencies. According to some such examples, the object renderer 910 may be capable of generating speaker feed signals in which the gain is adjusted (e.g., on a per-band basis), based on the characteristics of a  
10 particular speaker system.

**[00138]**       According to some implementations, the hearing environment data may include data regarding current environmental noise. For example, the apparatus 900 may receive a raw audio feed from a microphone or processed audio data that is based on audio signals from a microphone. In some implementations the apparatus 900 may include a  
15 microphone capable of providing data regarding current environmental noise. According to some such implementations, the object renderer 910 may be capable of generating speaker feed signals in which the gain is adjusted (e.g., on a per-band basis), based at least in part, on the current environmental noise. Additional examples of object renderer functionality are disclosed herein, including but not limited to the examples provided by Figure 11 and the  
20 corresponding description.

**[00139]**       The object renderer 910 may operate, at least in part, according to user input. In some examples, the object renderer 910 may be capable of modifying a distance (e.g., an angular separation) between at least some audio objects in the front area of a virtual acoustic space according to user input.

25           **[00140]**       Figure 9B shows an example of object prioritizers and object Tenderers in two different contexts. The apparatus 900a, which includes an object prioritizer 905a and an object renderer 910a, is capable of operating in a content creation context in this example. The content creation context may, for example, be an audio editing environment, such as a post-production editing environment, a sound effects creation environment, etc. Audio  
30 objects may be prioritized, e.g., by the object prioritizer 905a. According to some implementations, the object prioritizer 905a may be capable of determining suggested or default priority levels that a content creator could optionally adjust, according to user input. Corresponding audio object prioritization metadata may be created by the object prioritizer 905a and associated with audio objects.

[00141] The object renderer 910a may be capable of adjusting the levels of audio signals corresponding to audio objects according to the audio object prioritization metadata and of rendering the audio objects into a plurality of speaker feed signals based, at least in part, on the audio object position metadata. According to some such implementations, part of the content creation process may involve auditioning or testing the suggested or default priority levels determined by the object prioritizer 905a and adjusting the object prioritization metadata accordingly. Some such implementations may involve an iterative process of auditioning/testing and adjusting the priority levels according to a content creator's subjective impression of the audio playback, to ensure preservation of the content creator's creative intent. The object renderer 910a may be capable of adjusting audio object levels according to received hearing environment data. In some examples, the object renderer 910a may be capable of adjusting audio object levels according to a hearing loss model included in the hearing environment data.

[00142] The apparatus 900b, which includes an object prioritizer 905b and an object renderer 910b, is capable of operating in a consumer context in this example. The consumer context may, for example, be a cinema environment, a home theater environment, a mobile display device, etc. In this example, the apparatus 900b receives prioritization metadata along with audio objects, corresponding audio signals and other metadata, such as position metadata, size metadata, etc. In this implementation, the object renderer 910b produces speaker feed signals based on the audio object signals, audio object metadata and hearing environment data. In this example, the apparatus 900b includes an object prioritizer 905b, which may be convenient for instance in which the prioritization metadata is not available. In this implementation, the object prioritizer 905b is capable of making an audio object prioritization based, at least in part, on audio object type and of providing audio object prioritization metadata to the object renderer 910b. However, in alternative implementations the apparatus 900b may not include the object prioritizer 905b.

[00143] In the examples shown in Figure 9B, both the object prioritizer 905b and the object renderer 910b may optionally function according to received user input. For example, a consumer (such as a home theater owner or a cinema operator) may audition or "preview" rendered speaker feed signals according to one set of audio object prioritization metadata (e.g., a set of audio object prioritization metadata that is output from a content creation process) to determine whether the corresponding played-back audio is satisfactory for a particular reproduction environment. If not, in some implementations a user may invoke the operation of a local object prioritizer, such as the object prioritizer 905b, and may

optionally provide user input. This operation may produce a second set of audio object prioritization metadata that may be rendered into speaker feed signals by the object renderer 910b and auditioned. The process may continue until the consumer believes the resulting played-back audio is satisfactory.

5           **[00144]**       Figure 9C is a flow diagram that outlines one example of a method that may be performed by apparatus such as those shown in Figures, 7, 9A and/or 9B. For example, in some implementations method 950 may be performed by the apparatus 900a, in a content creation context. In some such examples, method 950 may be performed by the object prioritizer 905a. However, in some implementations method 950 may be performed  
10 by the apparatus 900b, in a consumer context. The blocks of method 950, like other methods described herein, are not necessarily performed in the order indicated. Moreover, such methods may include more or fewer blocks than shown and/or described.

**[00145]**       In this implementation, block 955 involves receiving audio data that includes a plurality of audio objects. In this example, the audio objects include audio signals  
15 and associated audio object metadata. In this implementation, the audio object metadata includes audio object position metadata. In some implementations, the audio object metadata may include one or more other types of audio object metadata, such as audio object type metadata, audio object size metadata, audio object prioritization metadata and/or one or more other types of audio object metadata.

20           **[00146]**       In the example shown in Figure 9C, block 960 involves extracting one or more features from the audio data. The features may, for example, include spectral flux, loudness, audio object size, entropy-related features, harmonicity features, spectral envelope features, phase features and/or temporal features. Some examples are described below with reference to Figure 10.

25           **[00147]**       In this implementation, block 965 involves determining an audio object type based, at least in part, on the one or more features extracted from the audio signals. In this example, the audio object type is selected from among a list of audio object types that includes dialogue. For example, a dialogue audio object may correspond to the speech of a particular individual. In some examples, the list of audio object types may  
30 include background music, events and/or ambiance. As noted above, in alternative implementations the audio object metadata may include audio object type metadata. According to some such implementations, determining the audio object type may involve evaluating the object type metadata.

**[00148]**       In the example shown in Figure 9C, block 970 involves making an

audio object prioritization based, at least in part, on the audio object type. In this example, the audio object prioritization determines, at least in part, a gain to be applied during a subsequent process of rendering the audio objects into speaker feed signals. In this implementation, making the audio object prioritization involves assigning a highest priority to audio objects that correspond to dialogue. In alternative implementations, making the audio object prioritization may involve assigning a highest priority to audio objects according to one or more other attributes, such as audio object volume or level. According to some implementations, some events (such as explosions, bullets sounds, etc.) may be assigned a higher priority than dialogue and other events (such as the sounds of a fire) may be assigned a lower priority than dialogue.

[00149] In this example, block 975 involves adding audio object prioritization metadata, based on the audio object prioritization, to the audio object metadata. The audio objects, including the corresponding audio signals and audio object metadata, may be provided to an audio object renderer.

[00150] As described in more detail below with reference to Figure 10, some implementations involve determining a confidence score regarding each audio object type determination and applying a weight to each confidence score to produce a weighted confidence score. The weight may correspond to the audio object type determination. In such implementations, making the audio object prioritization may be based, at least in part, on the weighted confidence score. In some examples, determining the audio object type may involve a machine learning method.

[00151] Some implementations of method 950 may involve receiving hearing environment data comprising a model of hearing loss and adjusting audio object levels according to the audio object prioritization and the hearing environment data. Such implementations also may involve rendering the audio objects into a plurality of speaker feed signals based, at least in part, on the audio object position metadata. Each speaker feed signal may correspond to at least one of the reproduction speakers within the reproduction environment.

[00152] The rendering process also may be based on reproduction environment data, which may include an express or implied indication of a number of reproduction speakers in a reproduction environment. In some examples, positions of reproduction speakers in the reproduction environment may be determined, or inferred, according to the reproduction environment configuration. Accordingly, the reproduction environment data may not need to include an express indication of positions of reproduction

speakers in the reproduction environment. In some implementations, the reproduction environment data may include an indication of a reproduction environment configuration.

[00153] In some examples, the reproduction environment data may include an indication of positions of reproduction speakers in the reproduction environment. In some implementations, the reproduction environment may be an actual reproduction environment, whereas in other implementations the reproduction environment may be a virtual reproduction environment. Accordingly, in some implementations the audio object position metadata may indicate locations in a virtual acoustic space.

[00154] As discussed elsewhere herein, in some instances the audio object metadata may include audio object size metadata. Some such implementations may involve receiving indications of a plurality of virtual speaker locations within the virtual acoustic space and rendering the audio objects to the plurality of virtual speaker locations within the virtual acoustic space based, at least in part, on the audio object position metadata and the audio object size metadata.

[00155] Figure 10 is a block diagram that shows examples of object prioritizer elements according to one implementation. The types and numbers of components shown in Figure 10 are merely shown by way of example. Alternative implementations may include more, fewer and/or different components. The object prioritizer 905c may, for example, be implemented via hardware, via software stored on non-transitory media, via firmware and/or by combinations thereof. In some examples, object prioritizer 905c may be implemented via a general purpose single- or multi-chip processor, a digital signal processor (DSP), an application specific integrated circuit (ASIC), a field programmable gate array (FPGA) or other programmable logic device, discrete gate or transistor logic, and/or discrete hardware components. In some implementations, the object prioritizer 905c may be implemented in an authoring/content creation context, such as an audio editing context for a video, for a movie, for a game, etc. However, in other implementations the object prioritizer 905c may be implemented in a cinema context, a home theater context, or another consumer-related context.

[00156] In this example, the object prioritizer 905c includes a feature extraction module 1005, which is shown receiving audio objects, including corresponding audio signals and audio object metadata. In this implementation, the feature extraction module 1005 is capable of extracting features from received audio objects, based on the audio signals and/or the audio object metadata. According to some such implementations, the set of features may be correspond with temporal, spectral and/or spatial properties of the

audio objects. The features extracted by the feature extraction module 1005 may include spectral flux, e.g., at the syllable rate, which may be useful for dialog detection. In some examples, the features extracted by the feature extraction module 1005 may include one or more entropy features, which may be useful for dialog and ambiance detection.

5           **[00157]**       According to some implementations, the features may include temporal features, one or more indicia of loudness and/or one or more indicia of audio object size, all of which may be useful for event detection. In some such implementations, the audio object metadata may include an indication of audio object size. In some examples, the feature extraction module 1005 may extract harmonicity features, which may be useful for  
10 dialog and background music detection. Alternatively, or additionally, the feature extraction module 1005 may extract spectral envelope features and/or phase features, which may be useful for modeling the spectral properties of the audio signals.

**[00158]**       In the example shown in Figure 10, the feature extraction module 1005 is capable of providing the extracted features 1007, which may include any  
15 combination of the above-mentioned features (and/or other features), to the classifier 1009. In this implementation, the classifier 1009 includes a dialogue detection module 1010 that is capable of detecting audio objects that correspond with dialogue, a background music detection module 1015 that is capable of detecting audio objects that correspond with background music, an event detection module 1020 that is capable of detecting audio objects  
20 that correspond with events (such as a bullet being fired, a door opening, an explosion, etc.) and an ambiance detection module 1025 that is capable of detecting audio objects that correspond with ambient sounds (such as rain, traffic sounds, wind, surf, etc.). In other examples, the classifier 1009 may include more or fewer elements.

**[00159]**       According to some examples, the classifier 1009 may be capable of  
25 implementing a machine learning method. For example, the classifier 1009 may be capable of implementing a Gaussian Mixture Model (GMM), a Support Vector Machine (SVM) or an Adaboost machine learning method. In some examples, the machine learning method may involve a "training" or set-up process. This process may have involved evaluating statistical properties of audio objects that are known to be particular audio object types, such  
30 as dialogue, background music, events, ambient sounds, etc. The modules of the classifier 1009 may have been trained to compare the characteristics of features extracted by the feature extraction module 1005 with "known" characteristics of such audio object types. The known characteristics may be characteristics of dialogue, background music, events, ambient sounds, etc., which have been identified by human beings and used as input for the

training process. Such known characteristics also may be referred to herein as "models."

**[00160]** In this implementation, the each element of the classifier 1009 is capable of generating and outputting a confidence score, which are shown as confidence scores 1030a-1030d in Figure 10. In this example, each of the confidence scores 1030a-  
5 1030d represents how close one or more characteristics of features extracted by the feature extraction module 1005 are to characteristics of a particular model. For example, if an audio object corresponds to people talking, then the dialog detection module 1010 may produce a high confidence score, whereas the background music detection module 1015 may produce a low confidence score.

10 **[00161]** In this example, the classifier 1009 is capable of applying a weighting factor  $W$  to each of the confidence scores 1030a-1030d. According to some implementations, each of the weighting factors  $W1-W4$  may be the results of a previous training process on manually labeled data, using a machine learning method such as one of those described above. In some implementations, the weighting factors  $W1-W4$  may have  
15 positive or negative constant values. In some examples, the weighting factors  $W1-W4$  may be updated from time to time according to relatively more recent machine learning results. The weighting factors  $W1-W4$  should result in priorities that provide improved experience for hearing impaired listeners. For example, in some implementations dialog may be assigned a higher priority, because dialogue is typically the most important part of an audio  
20 mix for a video, a movie, etc. Therefore, the weighting factor  $W1$  will generally be positive and larger than the weighting factors  $W2-W4$ . The resulting weighted confidence scores 1035a-1035d may be provided to the priority computation module 1050. If audio object prioritization metadata is available (for example, if audio object prioritization metadata is received with the audio signals and other audio object metadata, a weighting value  $Wp$  may  
25 be applied according to the audio object prioritization metadata.

**[00162]** In this example, the priority computation module 1050 is capable of calculating a sum of the weighted confidence scores in order to produce the final priority, which is indicated by the audio object prioritization metadata output by the priority computation module 1050. In alternative implementations, the priority computation module  
30 1050 may be capable of producing the final priority by applying one or more other types of functions to the weighted confidence scores. For example, in some instances the priority computation module 1050 may be capable of producing the final priority by applying a non-linear compressing function to the weighted confidence scores, in order to make the output within a predetermined range, for example between 0 and 1. If audio object prioritization

metadata is present for a particular audio object, the priority computation module 1050 may bias the final priority according to the priority indicated by the received audio object prioritization metadata. In this implementation, the priority computation module 1050 is capable of changing the priority assigned to audio objects according to optional user input.

5 For example, a user may be able to modify the weighting values in order to increase the priority of background music and/or another audio object type relative to dialogue, to increase the priority of a particular audio object as compared to the priority of other audio objects of the same audio object type, etc.

[00163] Figure 11 is a block diagram that shows examples of object renderer  
10 elements according to one implementation. The types and numbers of components shown in Figure 11 are merely shown by way of example. Alternative implementations may include more, fewer and/or different components. The object renderer 910c may, for example, be implemented via hardware, via software stored on non-transitory media, via firmware and/or by combinations thereof. In some examples, object renderer 910c may be implemented via a  
15 general purpose single- or multi-chip processor, a digital signal processor (DSP), an application specific integrated circuit (ASIC), a field programmable gate array (FPGA) or other programmable logic device, discrete gate or transistor logic, and/or discrete hardware components. In some implementations, the object renderer 910c may be implemented in an authoring/content creation context, such as an audio editing context for a video, for a movie,  
20 for a game, etc. However, in other implementations the object renderer 910c may be implemented in a cinema context, a home theater context, or another consumer-related context.

[00164] In this example, the object renderer 910c includes a leveling module  
1105, a pre-warping module 1110 and a rendering module 1115. In some implementations,  
25 as here, the rendering module 1115 includes an optional unmixer.

[00165] In this example, the leveling module 1105 is capable of receiving  
hearing environment data and of leveling audio signals based, at least in part, on the hearing  
environment data. In some implementations, the hearing environment data may include a  
hearing loss model, information regarding a reproduction environment (e.g., information  
30 regarding noise in the reproduction environment) and/or information regarding rendering  
hardware in the reproduction environment, such as information regarding the capabilities of  
one or more speakers of the reproduction environment. As noted above, the reproduction  
environment may include a headphone configuration, a virtual speaker configuration or an  
actual speaker configuration.

[00166] The leveling module 1105 may function in a variety of ways. The functionality of the leveling module 1105 may, for example, depend on the available processing power of the object renderer 910c. According to some implementations, the leveling module 1105 may operate according to a multiband compressor method. In some such implementations, adjusting the audio object levels may involve dynamic range compression. For example, referring to Figure 12, two DRC curves are shown. The solid line shows a sample DRC compression gain curve tuned for hearing loss. For example, audio objects with the highest priority may be leveled according to this curve. For a lower-priority audio object, the compression gain slopes may be adjusted as shown by the dashed line, receiving less boost and more cut than higher-priority audio objects. According to some implementations, the audio object priority may determine the degree of these adjustments. The units of the curves correspond to level or loudness. The DRC curves may be tuned per band according to, e.g., a hearing loss model, environmental noise, etc. Examples of more complex functionality of the leveling module 1105 are described below with reference to Figure 13. The output from the leveling module 1105 is provided to the rendering module 1115 in this example.

[00167] As described elsewhere herein, some implementations may involve increasing a distance between at least some audio objects in a virtual acoustic space. In some instances, the virtual acoustic space may include a front area and a back area. The front area and the back area may, for example, be determined relative to a position of a virtual listener's head in the virtual acoustic space.

[00168] According to some implementations, the pre-warping module 1110 may be capable of receiving audio object metadata, including audio object position metadata, and increasing a distance between at least some audio objects in the front area of the virtual acoustic space. Increasing this distance may, in some examples, improve the ability of a listener to hear the rendered audio objects more clearly. In some examples, the pre-warping module may adjust a distance between at least some audio objects according to user input. The output from the pre-warping module 1110 is provided to the rendering module 1115 in this example.

[00169] In this example, the rendering module 1115 is capable of rendering the output from the leveling module 1105 (and, optionally, output from the pre-warping module 1110) into speaker feed signals. As noted elsewhere, the speaker feed signals may correspond to virtual speakers or actual speakers.

[00170] In this example, the rendering module 1115 includes an optional

"unmixer." According to some implementations, the unmixer may apply special processing to at least some audio objects according to audio object size metadata. For example, the unmixer may be capable of determining whether an audio object has corresponding audio signals that include a directional component and a diffuse component. The unmixer may be capable of reducing a level of the diffuse component. According to some implementations, the unmixer may only apply such processing to audio objects that are at or above a threshold audio object size.

[00171] Figure 13 is a block diagram that illustrates examples of elements in a more detailed implementation. The types and numbers of components shown in Figure 13 are merely shown by way of example. Alternative implementations may include more, fewer and/or different components. The apparatus 900c may, for example, be implemented via hardware, via software stored on non-transitory media, via firmware and/or by combinations thereof. In some examples, apparatus 900c may be implemented via a general purpose single- or multi-chip processor, a digital signal processor (DSP), an application specific integrated circuit (ASIC), a field programmable gate array (FPGA) or other programmable logic device, discrete gate or transistor logic, and/or discrete hardware components. In some implementations, the apparatus 900c may be implemented in an authoring/content creation context, such as an audio editing context for a video, for a movie, for a game, etc. However, in other implementations the apparatus 900c may be implemented in a cinema context, a home theater context, or another consumer-related context.

[00172] In this example, the apparatus 900c includes a prioritizer 905d, excitation approximation modules 1325a-1325o, a gain solver 1330, an audio modification unit 1335 and a rendering unit 1340. In this particular implementation, the prioritizer 905d is capable of receiving audio objects 1 through N, prioritizing the audio objects 1 through N and providing corresponding audio object prioritization metadata to the gain solver 1330. In this example, the gain solver 1330 is a priority-weighted gain solver, which may function as described in the detailed discussion below.

[00173] In this implementation, the excitation approximation modules 1325b-1325o are capable of receiving the audio objects 1 through N, determining corresponding excitations  $E_1-E_N$  and providing the excitations  $E_1-E_N$  to the gain solver 1330. In this example, the excitation approximation modules 1325b-1325o are capable of determining corresponding excitations  $E_1-E_N$  based, in part, on the speaker deficiency data 1315a. The speaker deficiency data 1315a may, for example, correspond with linear frequency response deficiencies of one or more speakers. According to this example, the excitation

approximation module 1325a is capable of receiving environmental noise data 1310, determining a corresponding excitation  $E_0$  and providing the excitation  $E_0$  to the gain solver 1330.

[00174] In this example, other types of hearing environment data, including speaker deficiency data 1315b and hearing loss performance data 1320, are provided to the gain solver 1330. In this implementation, the speaker deficiency data 1315b correspond to speaker distortion. According to this implementation, the gain solver 1330 is capable of determining gain data based on the excitations  $E_0$   $E_N$  and the hearing environment data, and of providing the gain data to the audio modification unit 1335.

[00175] In this example, the audio modification unit 1335 is capable of receiving the audio objects 1 through N and modifying gains based, at least in part, on the gain data received from the gain solver 1330. Here, the audio modification unit 1335 is capable of providing gain-modified audio objects 1338 to the rendering unit 1340. In this implementation, the rendering unit 1340 is capable of generating speaker feed signals based on the gain-modified audio objects 1338.

[00176] The operation of the elements of Figure 13, according to some implementations, may be further understood with reference to the following remarks. Let  $LR_i$  be the loudness with which a person without hearing loss, in a noise-free playback environment, would perceive audio object  $i$ , after automatic gain control has been applied. This loudness, which may be calculated with a reference hearing model, depends on the level of all the other audio objects present. (In order to understand this phenomenon, consider that when another audio object is much louder than an audio object one may not be able to hear the audio object at all, so the audio object's perceived loudness is zero).

[00177] In general, loudness also depends on the environmental noise, hearing loss and speaker deficiencies. Under these conditions the same audio object will be perceived with a loudness  $LHL_i$ . This may be calculated using a hearing model  $H$  that includes hearing loss.

[00178] The goal of some implementations is to apply gains in spectral bands of each audio object such that, after these gains have been applied,  $LHL_i = LR_i$  for all the audio objects. In this case, every audio object may be perceived as a content creator intended for them to be perceived. If a person with reference hearing listened to the result, that person would perceive the result as if the audio objects had undergone dynamic range compression, as the signals inaudible to the person with hearing loss would have increased in

loudness and the signals that the person with hearing loss perceived as too loud would be reduced in loudness. This defines for us an objective goal of dynamic range compression matched to the environment.

[00179] However, because the perceived loudness of every audio object  
 5 depends on that of every other audio object, this cannot in general be satisfied for all audio objects. The solution of some disclosed implementations is to acknowledge that some audio objects may be more important than others, from a listener's point of view, and to assign audio object priorities accordingly. The *priority weighted gain solver* may be capable of calculating gains such that the difference or "distance" between  $LHL_i$  and  $LR_i$  is small for the  
 10 highest-priority audio objects and larger for lower-priority audio objects. This inherently results in reducing the gains on lower-priority audio objects, in order to reduce their influence. In one example, the gain solver calculates gains that minimize the following expression:

$$\min \sum_i p_i (LHL_i(b) - LR_i(b))^2 \quad (\text{Equation 2})$$

15 [00180] In Equation 2,  $p_i$  represents the priority assigned to audio object  $i$ , and  $LHL_i$  and  $LR_i$  are represented in the log domain. Other implementations may use other "distance" metrics, such as the absolute value of the loudness difference instead of the square of the loudness difference.

[00181] In one example of a hearing loss model, the loudness is calculated  
 20 from the sum of the specific loudness in each spectral band  $N(b)$ , which in turn is a function of the distribution of energy along the basilar membrane of the human ear, which we refer to herein as excitation  $E$ . Let  $E_i(b, ear)$  denote the excitation at the left or right ear due to audio object  $i$  in spectral band  $b$ .

$$LEU = \log \left[ \sum_{i, ears} N_{HL} \{ E_i(b, ear), \sum_{j \neq i} [g_j(b) E_j(b, ear)] + E_0(b, ear), \& \} \right] \quad (\text{Equation 3})$$

25 [00182] In Equation 3,  $E_0$  represents the excitation due to the environmental noise, and thus we will generally not be able to control gains for this excitation.  $N_{HL}$  represents the specific loudness, given the current hearing loss parameters, and  $g_i$  are the gains calculated by the priority-weighted gain solver for each audio object  $i$ . Similarly,

$$LR_i = \log \left[ \sum_{b, ears} N_R \{ E_i(b, ear), \sum_{j \neq i} [E_j(b, ear)] - E_0(b, ear), \& \} \right] \quad (\text{Equation 4})$$

[00183] In Equation 4,  $N_R$  is the specific loudness under the reference conditions of no hearing loss and no environmental noise. The loudness values and gains may undergo smoothing across time before being applied. Gains also may be smoothed across bands, to limit distortions caused by a filter bank.

5 [00184] In some implementations, the system may be simplified. For example, some implementations are capable of solving for a single broadband gain for the audio object  $i$  by letting  $g_{ib}=g^{\wedge}$ . Such implementations can, in some instances, dramatically reduce the complexity of the gain solver. Moreover, such implementations may improve the problem that users are used to listening through their hearing loss, and thus may sometimes  
10 be annoyed by the extra brightness if the highs are restored to the reference loudness levels.

[00185] Other simplified implementations may involve a hybrid system that has per-band gains for one type of audio object (e.g., for dialogue) and broadband gains for other types of audio object. Such implementations can making the dialog more intelligible, but can leave the timbre of the overall audio largely unchanged.

15 [00186] Still other simplified implementations may involve making some assumptions regarding the spectral energy distribution, e.g., by measuring in fewer bands and interpolating. One very simple approach involves making an assumption that all audio objects are pink noise signals and simply measuring the audio objects' energy. Then,  $E.(b)=k^{\wedge}E.$ , where the constants  $k_b$  represent the pink noise characteristic.

20 [00187] The passage of the audio through the speaker's response, and through the head-related transfer function, then through the outer ear and the middle ear can all be modeled with a relatively simple transfer function. An example of a middle ear transfer function is shown in Fig.1 of Brian C. J. Moore and Brian R. Glasberg, *A Revised Model of Loudness Perception Applied to Cochlear Hearing Loss*, in *Hearing Research*, 188 (2004)  
25 ("MG2004") and the corresponding discussion, which are hereby incorporated by reference. An example transfer function through the outer ear for a frontal source is given in Figure 1 of Brian R. Glasberg, Brian C. J. Moore and Thomas Baer, *A Model for the Prediction of Thresholds, Loudness and Partial Loudness*, in *J. Audio Eng. Soc*, 45 (1997) ("MG1997") and the corresponding discussion, which are hereby incorporated by reference.

30 [00188] The spectrum of each audio object is banded into equivalent rectangular bands *ERB* that model the logarithmic spacing with frequency along the basilar membrane via a filterbank, and the energy out of these filters is smoothed to give the

excitation  $E_i(b)$ , where  $b$  indexes over the *ERB*.

[00189] In some examples, the nonlinear compression from excitation  $E_i(b)$  to specific loudness  $N_i(b)$  may be calculated via the following equation:

$$N_i(b) = C [ [G(b)E_i(b) + A(b)]^f - A(b) ]^{\alpha_i}. \quad \text{(Equation 5)}$$

5 [00190] In Equation 5,  $A$ ,  $G$  and  $\alpha$  represent an interdependent function of  $b$ . For example,  $G$  may be matched to experimental data and then the corresponding  $\alpha$  and  $A$  values can be calculated. Some examples are provided in figures 2, 3 and 4 of MG2004 and the corresponding discussion, which are hereby incorporated by reference.

[00191] The value at levels close to the absolute threshold of hearing in a quiet environment actually falls off more quickly than Equation 5 suggests. (It is not necessarily zero because even though a tone at a single frequency may be inaudible if a sound is wideband, the combination of inaudible tones can be audible.) A correction factor of  $[2E_i(b) / (E_i(b) + E_{THRQ}(b))]^{1.5}$  may be applied when  $E_i(b) < E_{THRQ}(b)$ . An example of  $E_{mRQ}(b)$  is given in Fig 2 of MG2004 and  $G \cdot E_{RQ} = const$ . The model can also be made more accurate at very high levels of excitation  $E_i > 10^{10}$  by using the alternative equation  $N_i(b) = C(E_i(b)/1.115)$ .

[00192] The perceived reference loudness may then be calculated by summing the specific loudness over bands and ears, e.g., according to the following:

$$LR_i = 10 \log_{10} [ \sum_b (N_i(left\ ear, b) + N_i(right\ ear, b)) ] \quad \text{(Equation 6)}$$

20 [00193] The reference loudness model is not yet complete because we should also incorporate the effects of the other audio objects present at the time. Even though environmental noise is not included for the reference model, the loudness of audio object  $i$  is calculated in the presence of audio objects  $j \neq i$ . This is discussed in the next section.

[00194] Consider sounds with intermediate range audio, wherein  $E_{THRQ} < (E_i + E_n) < 10^{10}$ , where we use  $E_n = \sum_{j \neq i} g_j(b)E_j + E_0$  to represent the excitation from all other background audio objects including environmental noise, and the absolute threshold noise may be represented as  $E_{THRQ} = K E_{noise} + E_{THRQ}$ , where  $K$  is a function of frequency found in Fig 9 of MG1997 and the corresponding discussion, which is hereby incorporated by reference. In this situation, the total specific loudness  $N_i$  will be

$$N_t = C[(E_i + E_n)G + A]^\alpha - A^\alpha \quad (\text{Equation 7})$$

**[00195]** Some of this loudness will be perceived as coming from the audio object  $i$ . It will be assumed that  $N_t = N_i + N_n - N_i$  may be partitioned between  $N_i$  and  $N_n$ , e.g., as described in the examples provided in MG1997 (which are hereby incorporated by

5

$$N_i = C[(E_i + E_n)G + A]^\alpha - A^\alpha -$$

$$C \left\{ [(E_n(1+K) + E_{THRQ})G + A]^\alpha - (E_{THRQ}G + A)^\alpha \right\} \left( \frac{E_{THRQ}}{E_i} \right)^{0.3}$$

**[00196]** When  $E_i < E_{THRQ}$  we may use an alternative equation:

$$N_i = C \left( \frac{2E_i}{E_i + E_n} \right)^{1.5} \left\{ \frac{(E_{THRQ}G + A)^\alpha - A^\alpha}{[(E_n(1+K) + E_{THRQ})G + A]^\alpha - (E_nG + A)^\alpha} \right\}$$

$$10 \cdot \left\{ [(E_i + E_n)G + A]^\alpha - (E_nG + A)^\alpha \right\}$$

**[00197]** More accurate equations for situations in which  $E_i + E_n > 10^{10}$ , such as equations 19 and 20 of MG1997, which are hereby incorporated by reference.

**[00198]** The effects of hearing loss may include: (1) an elevation of the absolute threshold in quiet; (2) a reduction in (or loss of) the compressive non-linearity; (3) a loss of frequency selectivity and/or (4) "dead regions" in the cochlea, with no response at all. Some implementations disclosed herein address the first two effects by fitting a new value of  $G(b)$  to the hearing loss and recalculating the corresponding  $a$  and  $A$  values for this value. Some such methods may involve adding an attenuation to the excitation that may scale with the level above absolute threshold. Some implementations disclosed herein address the third effect by fitting a broadening factor to the calculation of the spectral bands, which in some implementations may be equivalent rectangular (ERB) bands. Some implementations disclosed herein address the fourth effect by setting  $g_i(b) = 0$ , and  $E_i(b) = 0$  for the dead region. After the priority-based gain solver has then solved for the other gains with  $g_i(b) = 0$ , this  $g_i(b)$  may be replaced with a value that is close enough to its neighbors  $g_i(b+1)$  and  $g_i(b-1)$ ,

20 so that the value does not cause distortions in the filterbank.

**[00199]** According to some examples, the foregoing effects of hearing loss may be addressed by extrapolating from an audiogram and assuming that the total hearing

loss is divided into outer hearing loss and inner hearing loss. Some relevant examples are described in MG2004 and are hereby incorporated by reference. For example, Section 3.1 of MG2004 explains that one may obtain the total hearing loss for each band by interpolating this from the audiogram. A remaining problem is to distinguish outer hearing loss (OHL) from inner hearing loss (IHL). Setting  $OHL = 0.9 \text{ THL}$  provides a good solution to this problem in many instances.

**[00200]** Some implementations involve compensating for speaker deficiencies by applying a speaker transfer function to the  $E_n$  when calculating  $LHL$ . In practice, however, below a certain frequency a speaker generally produces little energy and creates significant distortion. For these frequencies, some implementations involve setting  $E_n(b)=0$  and  $g_i(b)=0$ .

**[00201]** Various modifications to the implementations described in this disclosure may be readily apparent to those having ordinary skill in the art. The general principles defined herein may be applied to other implementations without departing from the scope of this disclosure. Thus, the claims are not intended to be limited to the implementations shown herein, but are to be accorded the widest scope consistent with this disclosure, the principles and the novel features disclosed herein.

**CLAIMS**

1. A method, comprising:
  - receiving audio data comprising a plurality of audio objects, the audio objects including audio signals and associated audio object metadata, the audio object metadata
  - 5 including audio object position metadata;
  - receiving reproduction environment data comprising an indication of a number of reproduction speakers in a reproduction environment;
  - determining at least one audio object type from among a list of audio object types that includes dialogue;
  - 10 making an audio object prioritization based, at least in part, on the audio object type, wherein making the audio object prioritization involves assigning a highest priority to audio objects that correspond to dialogue;
  - adjusting audio object levels according to the audio object prioritization; and
  - rendering the audio objects into a plurality of speaker feed signals based, at least in
  - 15 part, on the audio object position metadata, wherein each speaker feed signal corresponds to at least one of the reproduction speakers within the reproduction environment.
2. The method of claim 1, further comprising receiving hearing environment data comprising at least one factor selected from a group of factors consisting of: a model of hearing loss; a deficiency of at least one reproduction speaker; and current environmental
- 20 noise, wherein adjusting the audio object levels is based, at least in part, on the hearing environment data.
3. The method of claim 1 or claim 2, wherein rendering involves rendering the audio objects to locations in a virtual acoustic space.
4. The method of claim 3, wherein the rendering involves increasing a distance between
- 25 at least some audio objects in the virtual acoustic space.
5. The method of claim 3 or claim 4, wherein the virtual acoustic space includes a front area and a back area and wherein the rendering involves increasing a distance between at least some audio objects in the front area of the virtual acoustic space.

6. The method of any one of claims 3-5, wherein the rendering involves rendering the audio objects according to a plurality of virtual speaker locations within the virtual acoustic space.
7. The method of any one of claims 1-6, wherein the audio object metadata includes  
5 metadata indicating audio object size and wherein making the audio object prioritization involves applying a function that reduces a priority of non-dialogue audio objects according to increases in audio object size.
8. The method of any one of claims 1-7, further comprising selecting at least one audio object that will not be rendered based, at least in part, on the audio object prioritization.
- 10 9. The method of any one of claims 1-8, wherein adjusting the audio object levels involves differentially adjusting levels in frequency bands of corresponding audio signals.
10. The method of any one of claims 1-9, wherein the audio object metadata includes audio object prioritization metadata and wherein adjusting the audio object levels is based, at least in part, on the audio object prioritization metadata.
- 15 11. The method of any one of claims 1-10, further comprising:  
determining that an audio object has audio signals that include a directional component and a diffuse component; and  
reducing a level of the diffuse component.
12. The method of any one of claims 1-11, wherein adjusting the audio object levels  
20 involves dynamic range compression.
13. The method of any one of claims 1-12, wherein the reproduction environment comprises a Dolby Surround 5.1 configuration, a Dolby Surround 7.1 configuration, a Hamasaki 22.2 surround sound configuration, a headphone configuration, a Dolby Surround 5.1.2 configuration, a Dolby Surround 7.1.2 configuration or a Dolby Atmos configuration.
- 25 14. The method of any one of claims 1-13, wherein the audio object metadata includes audio object type metadata and wherein determining the audio object type involves evaluating the object type metadata.

15. The method of any one of claims 1-13, wherein determining the audio object type involves analyzing the audio signals of audio objects.
16. The method of any one of claims 1-15, wherein the list of audio object types also includes background music, events and ambiance.
- 5 17. A method, comprising:  
receiving audio data comprising a plurality of audio objects, the audio objects including audio signals and associated audio object metadata;  
extracting one or more features from the audio data;  
determining an audio object type based, at least in part, on features extracted from the  
10 audio signals, wherein the audio object type is selected from a list of audio object types that includes dialogue;  
making an audio object prioritization based, at least in part, on the audio object type, wherein the audio object prioritization determines, at least in part, a gain to be applied during a process of rendering the audio objects into speaker feed signals and wherein making the  
15 audio object prioritization involves assigning a highest priority to audio objects that correspond to dialogue; and  
adding audio object prioritization metadata, based on the audio object prioritization, to the audio object metadata.
18. The method of claim 17, wherein the list of audio object types also includes  
20 background music, events and ambiance.
19. The method of claim 17, wherein the one or more features include at least one feature from a list of features consisting of: spectral flux; loudness; audio object size; entropy-related features; harmonicity features; spectral envelope features; phase features; and temporal features.
- 25 20. The method of any one of claims 17-19, further comprising:  
determining a confidence score regarding each audio object type determination; and  
applying a weight to each confidence score to produce a weighted confidence score, the weight corresponding to the audio object type determination, wherein making an audio object prioritization is based, at least in part, on the weighted confidence score.

21. The method of any one of claims 17-20, wherein determining the audio object type involves a machine learning method.
22. The method of any one of claims 17-21, further comprising:  
receiving hearing environment data comprising a model of hearing loss;  
5 adjusting audio object levels according to the audio object prioritization and the hearing environment data; and  
rendering the audio objects into a plurality of speaker feed signals based, at least in part, on the audio object position metadata, wherein each speaker feed signal corresponds to at least one of the reproduction speakers within the reproduction environment.
- 10 23. The method of any one of claims 17-21, wherein the audio object metadata includes audio object size metadata and wherein the audio object position metadata indicates locations in a virtual acoustic space, further comprising:  
receiving hearing environment data comprising a model of hearing loss;  
receiving indications of a plurality of virtual speaker locations within the virtual  
15 acoustic space;  
adjusting audio object levels according to the audio object prioritization and the hearing environment data; and  
rendering the audio objects to the plurality of virtual speaker locations within the virtual acoustic space based, at least in part, on the audio object position metadata and the  
20 audio object size metadata.
24. An apparatus, comprising:  
means for receiving audio data comprising a plurality of audio objects, the audio objects including audio signals and associated audio object metadata, the audio object metadata including at least audio object position metadata;  
25 means for receiving reproduction environment data comprising an indication of a number of reproduction speakers in a reproduction environment;  
means for determining at least one audio object type from among a list of audio object types that includes dialogue;  
means for making an audio object prioritization based, at least in part, on the audio  
30 object type, wherein making the audio object prioritization involves assigning a highest priority to audio objects that correspond to dialogue;

means for adjusting audio object levels according to the audio object prioritization;  
and

means for rendering the audio objects into a plurality of speaker feed signals based, at least in part, on the audio object position metadata, wherein each speaker feed signal  
5 corresponds to at least one of the reproduction speakers within the reproduction environment.

25. The apparatus of claim 24, wherein the means for receiving audio data includes means for receiving hearing environment data comprising at least one factor selected from a group of factors consisting of: a model of hearing loss; a deficiency of at least one reproduction speaker; and current environmental noise, wherein adjusting the audio object levels is based,  
10 at least in part, on the hearing environment data.

26. An apparatus, comprising:

an interface system capable of receiving audio data comprising a plurality of audio objects, the audio objects including audio signals and associated audio object metadata, the audio object metadata including at least audio object position metadata; and

15 a control system capable of:

receiving reproduction environment data comprising an indication of a number of reproduction speakers in a reproduction environment;

determining at least one audio object type from among a list of audio object types that includes dialogue;

20 making an audio object prioritization based, at least in part, on the audio object type, wherein making the audio object prioritization involves assigning a highest priority to audio objects that correspond to dialogue;

adjusting audio object levels according to the audio object prioritization; and

25 rendering the audio objects into a plurality of speaker feed signals based, at least in part, on the audio object position metadata, wherein each speaker feed signal corresponds to at least one of the reproduction speakers within the reproduction environment.

27. The apparatus of claim 26, wherein the interface system is capable of receiving hearing environment data comprising at least one factor selected from a group of factors  
30 consisting of: a model of hearing loss; a deficiency of at least one reproduction speaker; and current environmental noise, and wherein the control system is capable of adjusting the audio object levels based, at least in part, on the hearing environment data.

28. A non-transitory medium having software stored thereon, the software including instructions for controlling at least one device for:

receiving audio data comprising a plurality of audio objects, the audio objects including audio signals and associated audio object metadata, the audio object metadata including at least audio object position metadata;

receiving reproduction environment data comprising an indication of a number of reproduction speakers in a reproduction environment;

determining at least one audio object type from among a list of audio object types that includes dialogue;

making an audio object prioritization based, at least in part, on the audio object type, wherein making the audio object prioritization involves assigning a highest priority to audio objects that correspond to dialogue;

adjusting audio object levels according to the audio object prioritization; and

rendering the audio objects into a plurality of speaker feed signals based, at least in part, on the audio object position metadata, wherein each speaker feed signal corresponds to at least one of the reproduction speakers within the reproduction environment.

29. The non-transitory medium of claim 28, wherein the software includes instructions for controlling the at least one device to receive hearing environment data comprising at least one factor selected from a group of factors consisting of: a model of hearing loss; a deficiency of at least one reproduction speaker; and current environmental noise, and wherein adjusting the audio object levels is based, at least in part, on the hearing environment data.

30. An apparatus, comprising:

means for receiving audio data comprising a plurality of audio objects, the audio objects including audio signals and associated audio object metadata;

means for extracting one or more features from the audio data;

means for determining an audio object type based, at least in part, on features extracted from the audio signals, wherein the audio object type is selected from a list of audio object types that includes dialogue;

means for making an audio object prioritization based, at least in part, on the audio object type, wherein the audio object prioritization determines, at least in part, a gain to be applied during a process of rendering the audio objects into speaker feed signals and wherein

making the audio object prioritization involves assigning a highest priority to audio objects that correspond to dialogue; and

means for adding audio object prioritization metadata, based on the audio object prioritization, to the audio object metadata.

5 31. The apparatus of claim 30, further comprising:

means for receiving hearing environment data comprising a model of hearing loss;

means for adjusting audio object levels according to the audio object prioritization and the hearing environment data; and

10 means for rendering the audio objects into a plurality of speaker feed signals based, at least in part, on the audio object position metadata, wherein each speaker feed signal corresponds to at least one of the reproduction speakers within the reproduction environment.

32. An apparatus, comprising:

an interface system capable of receiving audio data comprising a plurality of audio objects, the audio objects including audio signals and associated audio object metadata; and  
15 a control system capable of:

extracting one or more features from the audio data;

determining an audio object type based, at least in part, on features extracted from the audio signals, wherein the audio object type is selected from a list of audio object types that includes dialogue;

20 making an audio object prioritization based, at least in part, on the audio object type, wherein the audio object prioritization determines, at least in part, a gain to be applied during a process of rendering the audio objects into speaker feed signals and wherein making the audio object prioritization involves assigning a highest priority to audio objects that correspond to dialogue; and

25 adding audio object prioritization metadata, based on the audio object prioritization, to the audio object metadata.

33. The apparatus of claim 32, wherein the interface system capable of receiving hearing environment data comprising a model of hearing loss and wherein the control system is capable of:

30 adjusting audio object levels according to the audio object prioritization and the hearing environment data; and

rendering the audio objects into a plurality of speaker feed signals based, at least in part, on the audio object position metadata, wherein each speaker feed signal corresponds to at least one of the reproduction speakers within the reproduction environment.

34. The apparatus of claim 32, wherein the interface system capable of receiving hearing environment data and wherein the control system comprises at least one excitation approximation module capable of determining excitation data that include an excitation for each of the plurality of audio objects, wherein an excitation is a function of a distribution of energy along a basilar membrane of a human ear and wherein at least one of the excitations is based, at least in part, on hearing environment data.

35. The apparatus of claim 34, wherein the control system includes a gain solver capable of receiving the excitation data and of determining gain data based, at least in part, on the excitations, the audio object prioritization and the hearing environment data.

36. The apparatus of claim 34 or 35, wherein the hearing environment data comprises at least one type of data selected from a group of data types consisting of environmental noise data, speaker deficiency data and hearing loss performance data.

37. A non-transitory medium having software stored thereon, the software including instructions for controlling at least one device for:

receiving audio data comprising a plurality of audio objects, the audio objects including audio signals and associated audio object metadata;

extracting one or more features from the audio data;

determining an audio object type based, at least in part, on features extracted from the audio signals, wherein the audio object type is selected from a list of audio object types that includes dialogue;

making an audio object prioritization based, at least in part, on the audio object type, wherein the audio object prioritization determines, at least in part, a gain to be applied during a process of rendering the audio objects into speaker feed signals and wherein making the audio object prioritization involves assigning a highest priority to audio objects that correspond to dialogue; and

adding audio object prioritization metadata, based on the audio object prioritization, to the audio object metadata.

38. The non-transitory medium of claim 37, wherein the software includes instructions for controlling the at least one device for:

receiving hearing environment data comprising a model of hearing loss;

adjusting audio object levels according to the audio object prioritization and the

5 hearing environment data; and

rendering the audio objects into a plurality of speaker feed signals based, at least in part, on the audio object position metadata, wherein each speaker feed signal corresponds to at least one of the reproduction speakers within the reproduction environment.

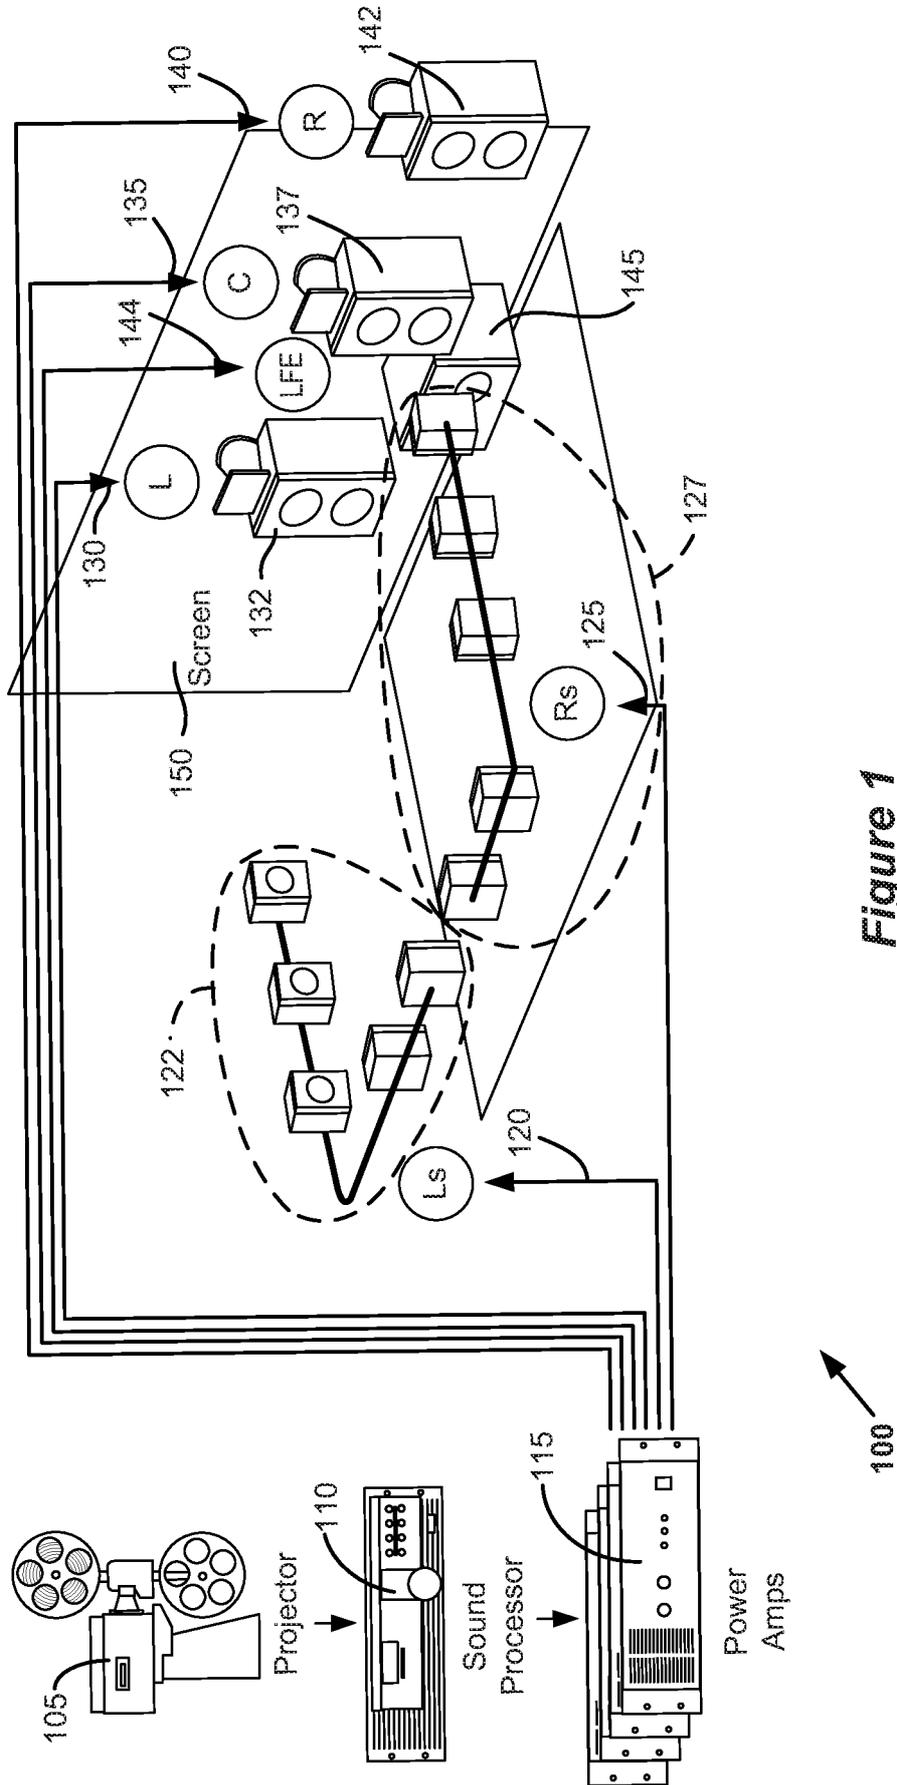


Figure 1

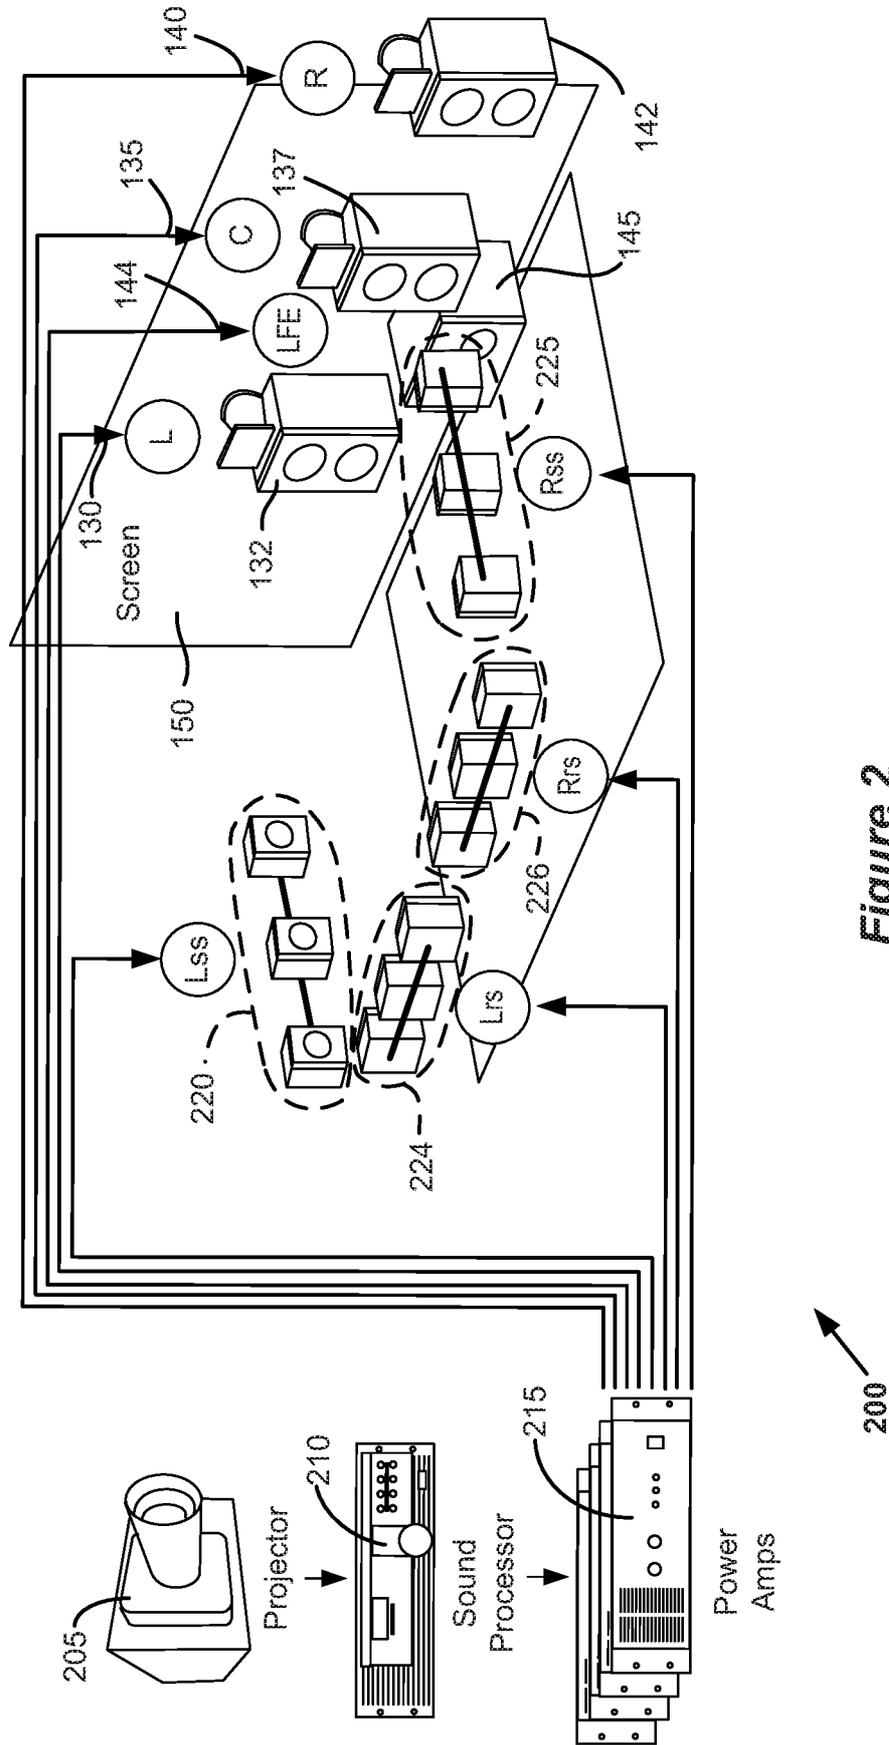


Figure 2

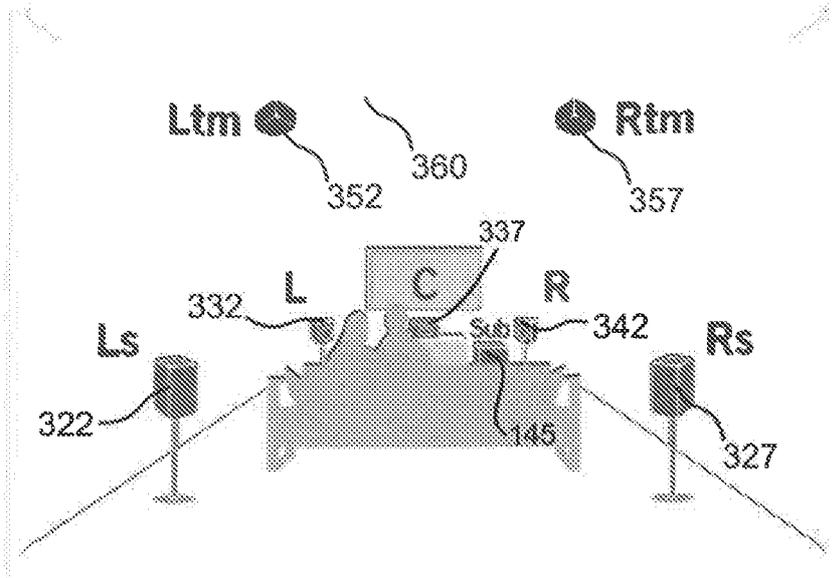


Figure 3A

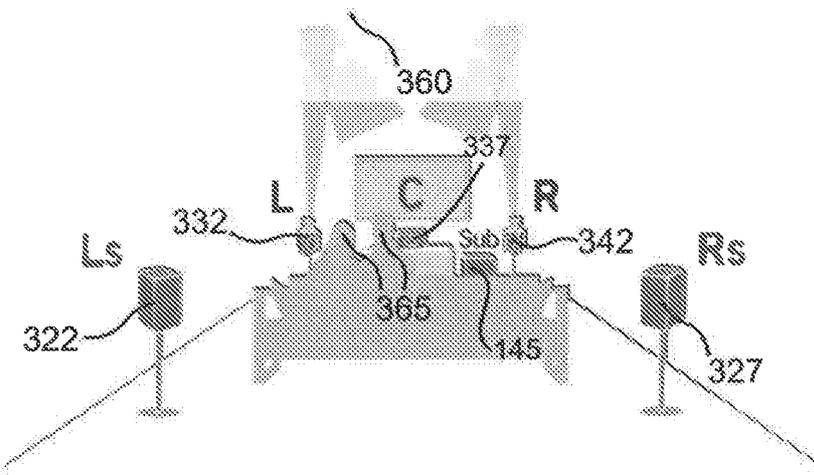


Figure 3B

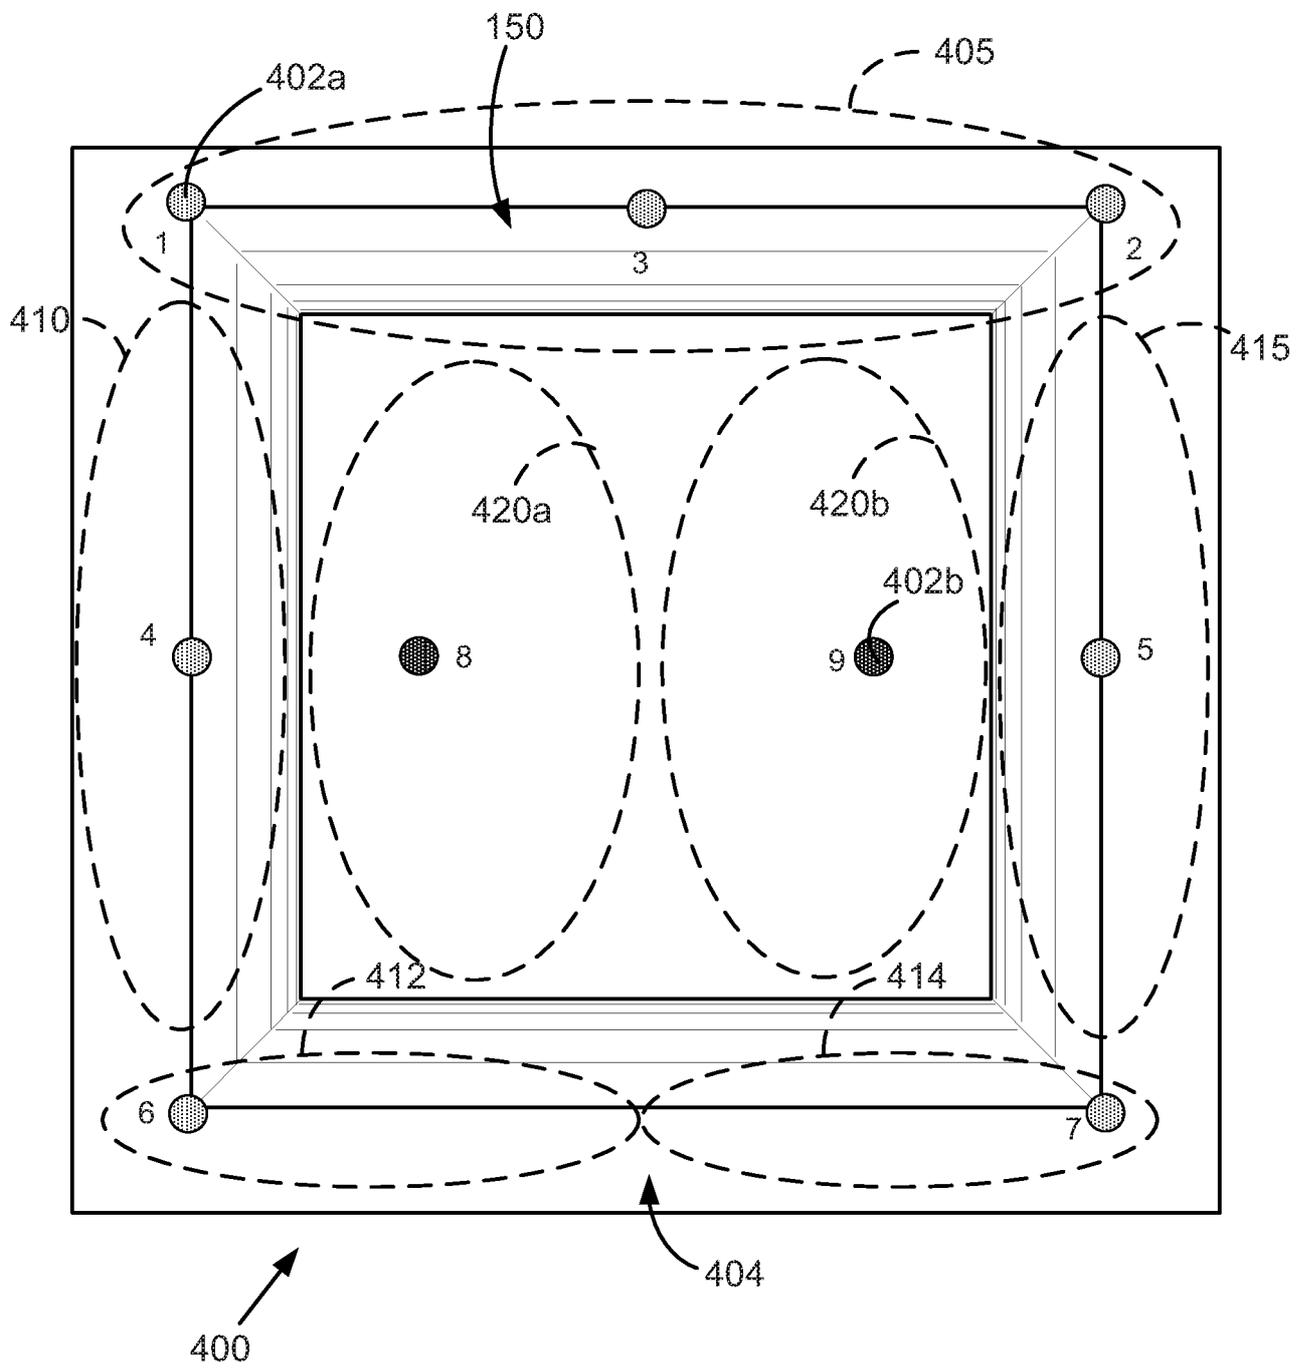


Figure 4A

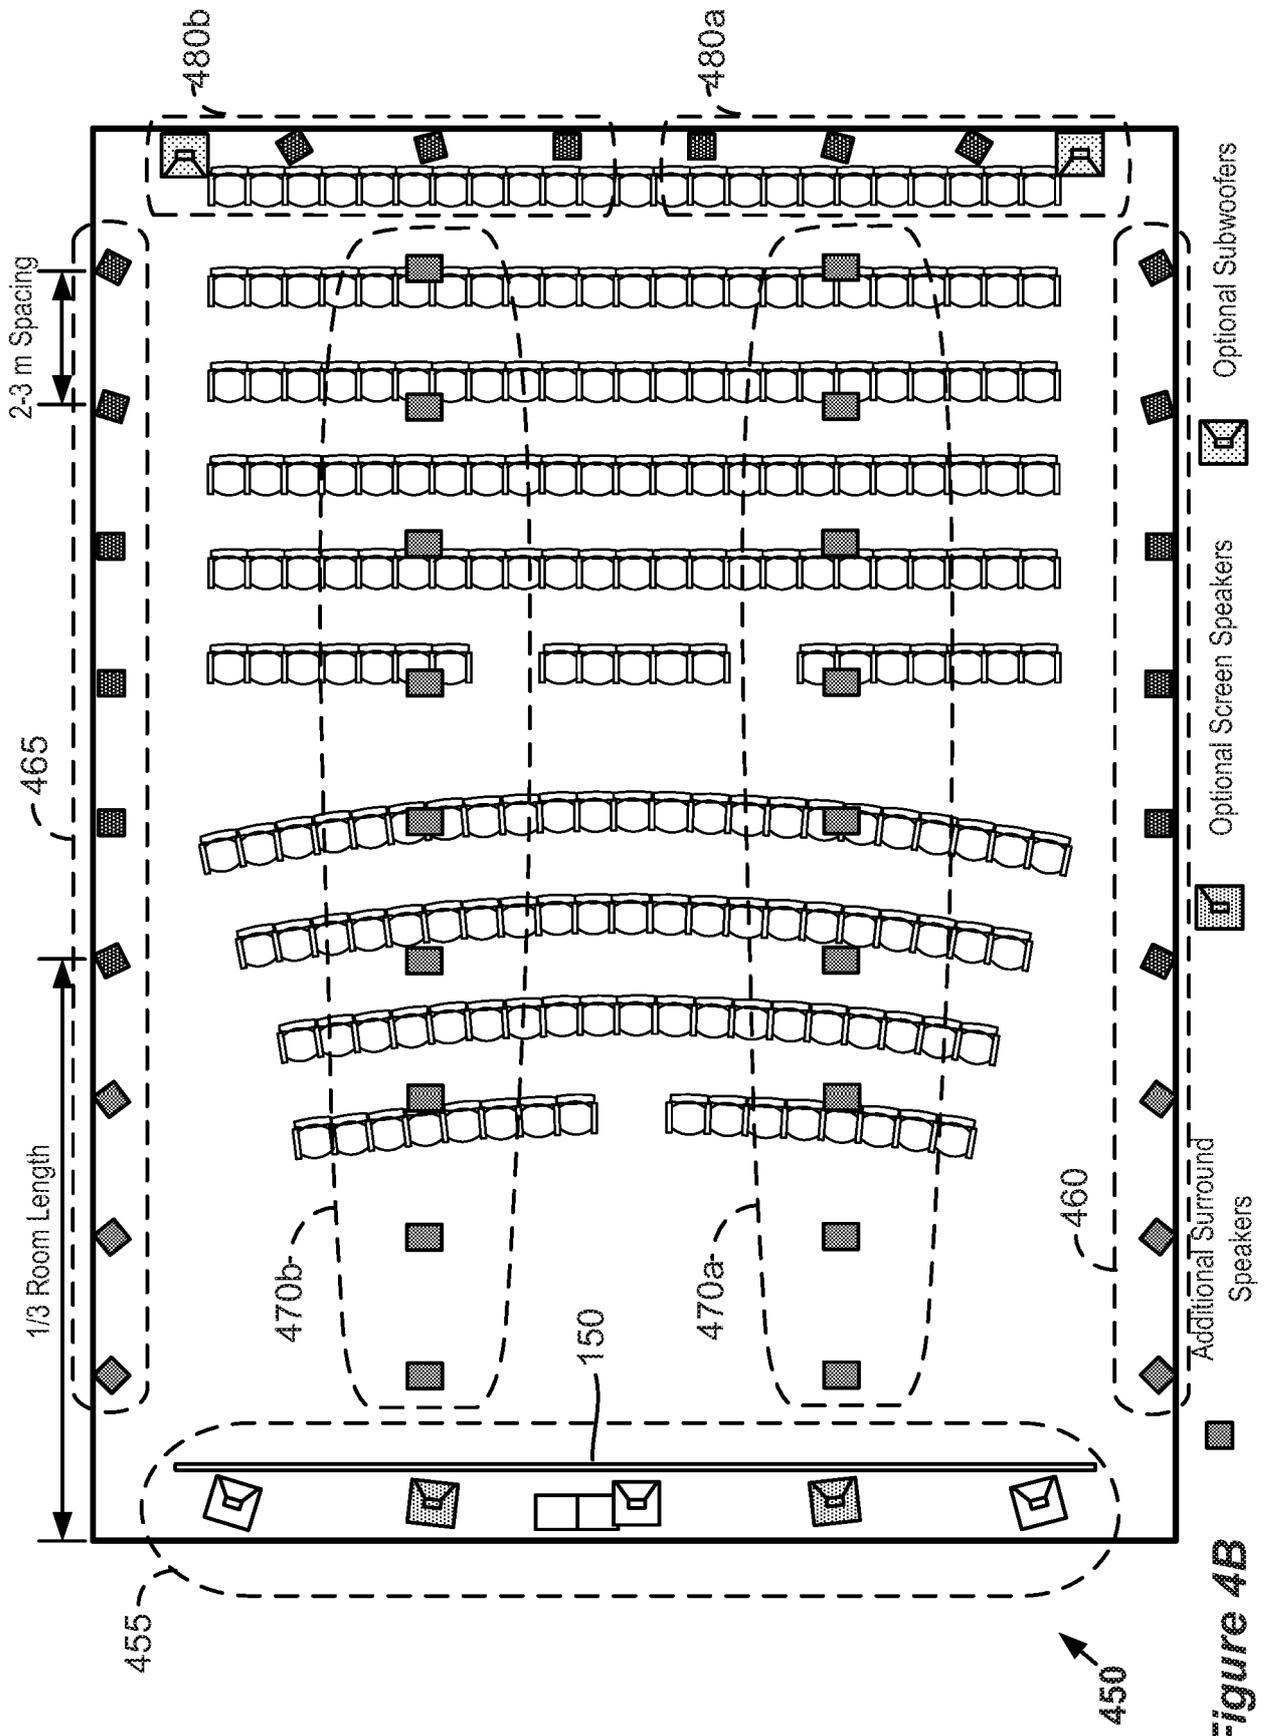


Figure 4B

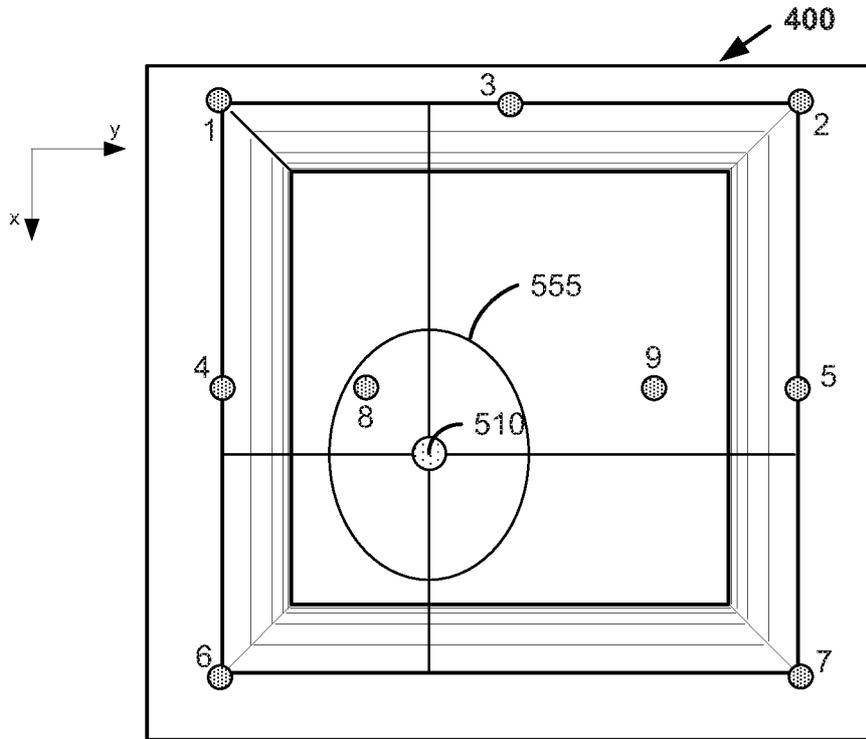


Figure 5A

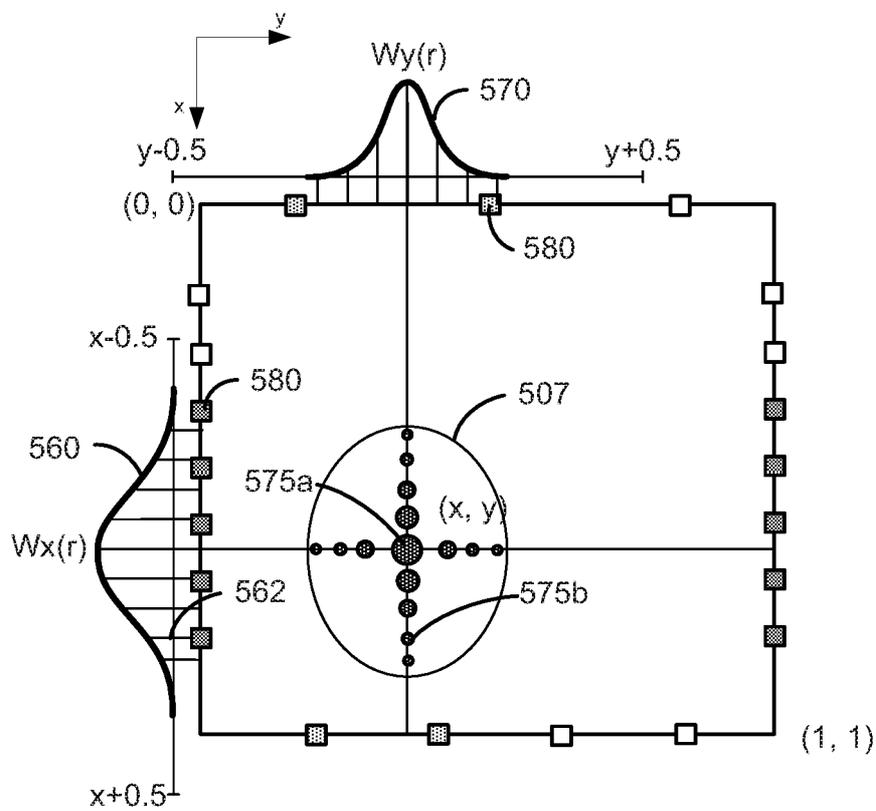
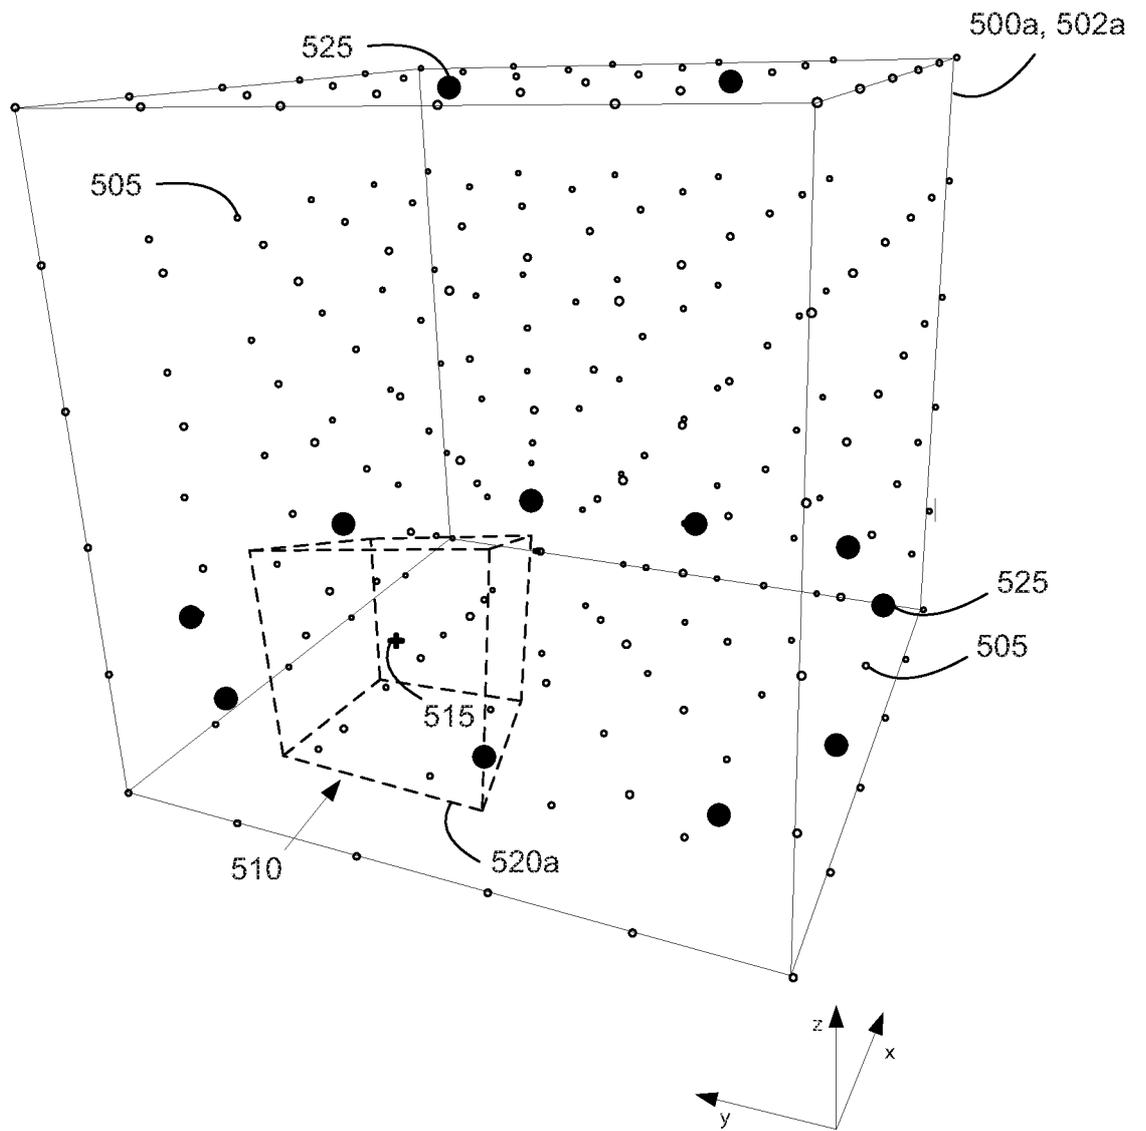
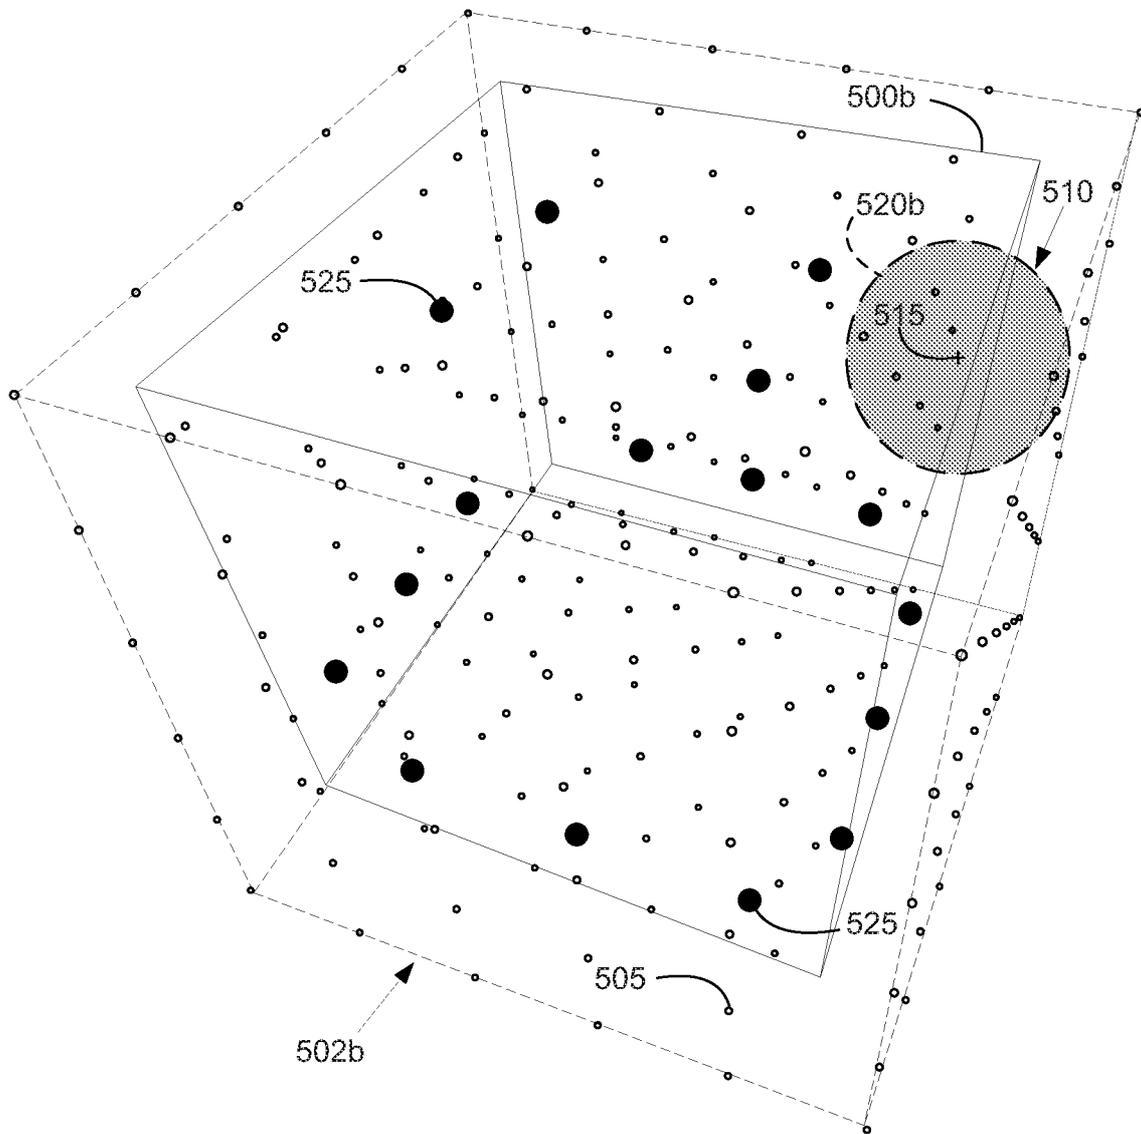


Figure 5B



**Figure 5C**



**Figure 5D**

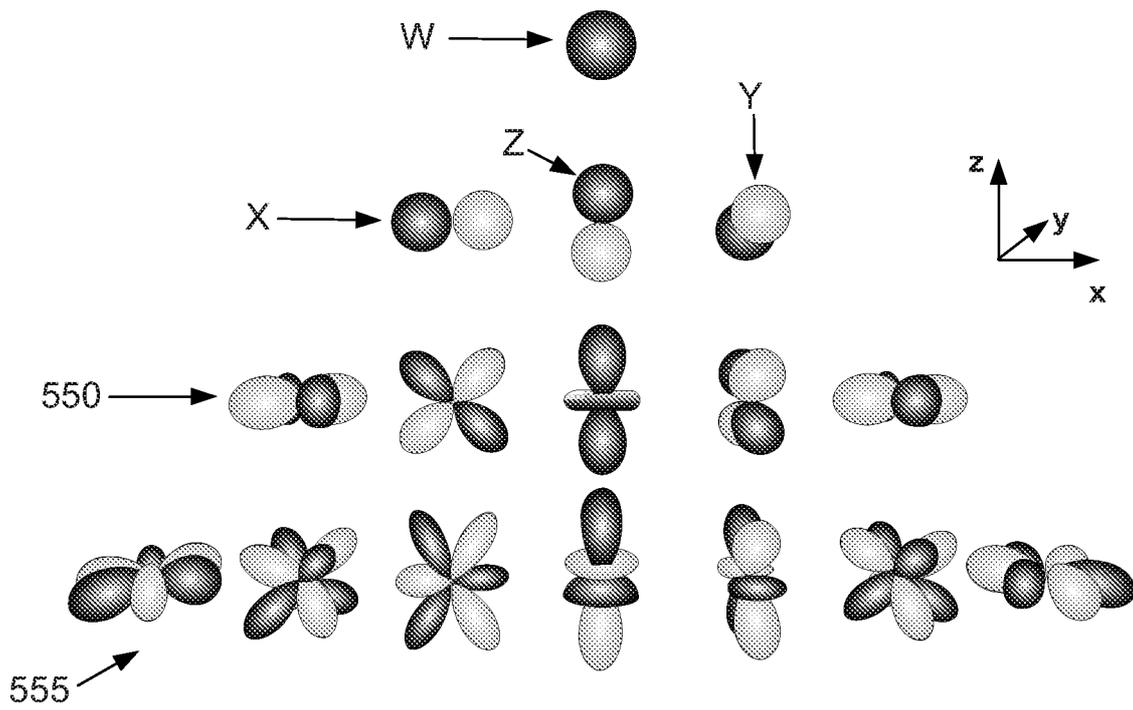
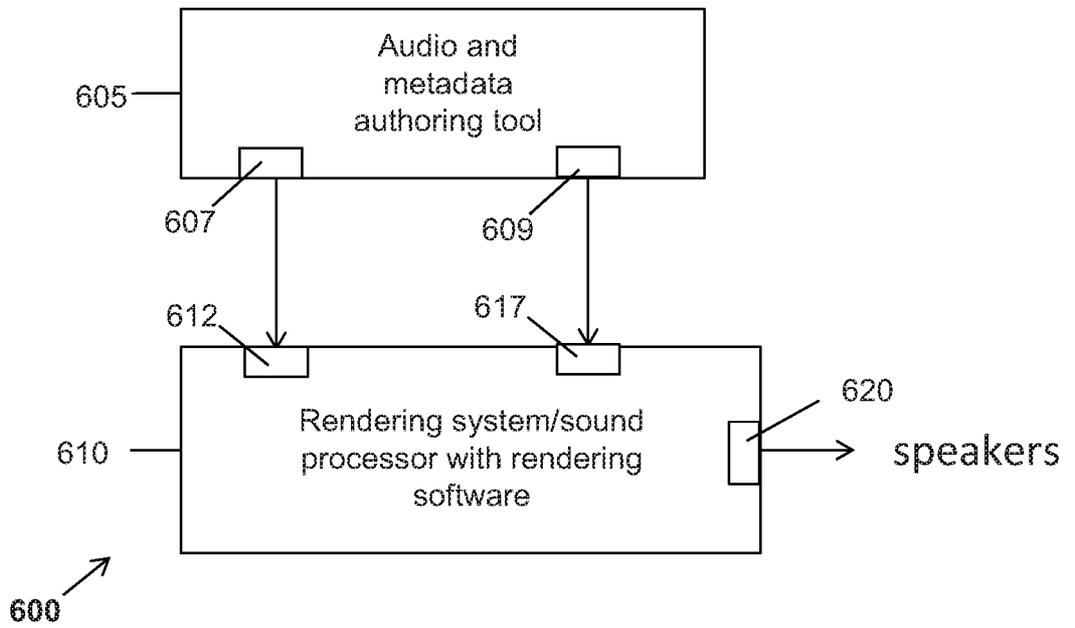
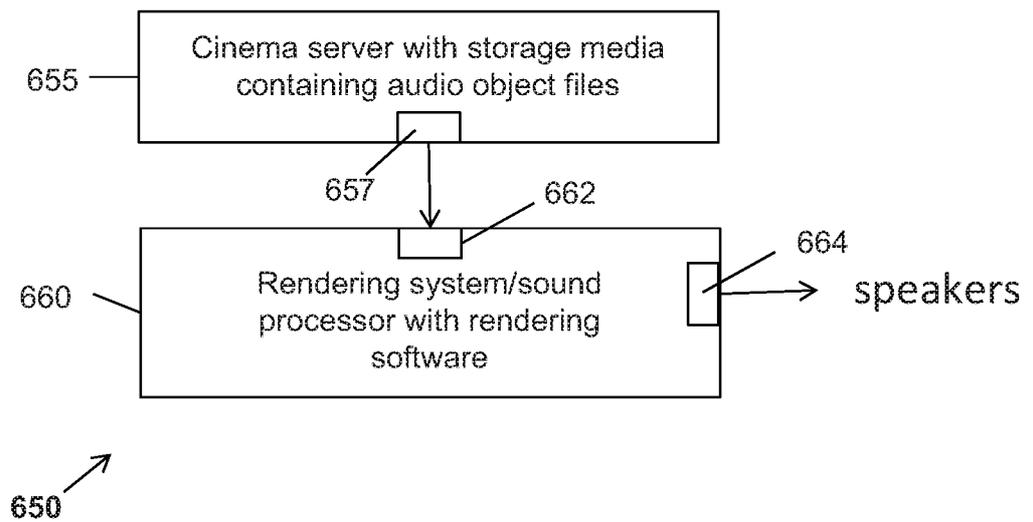


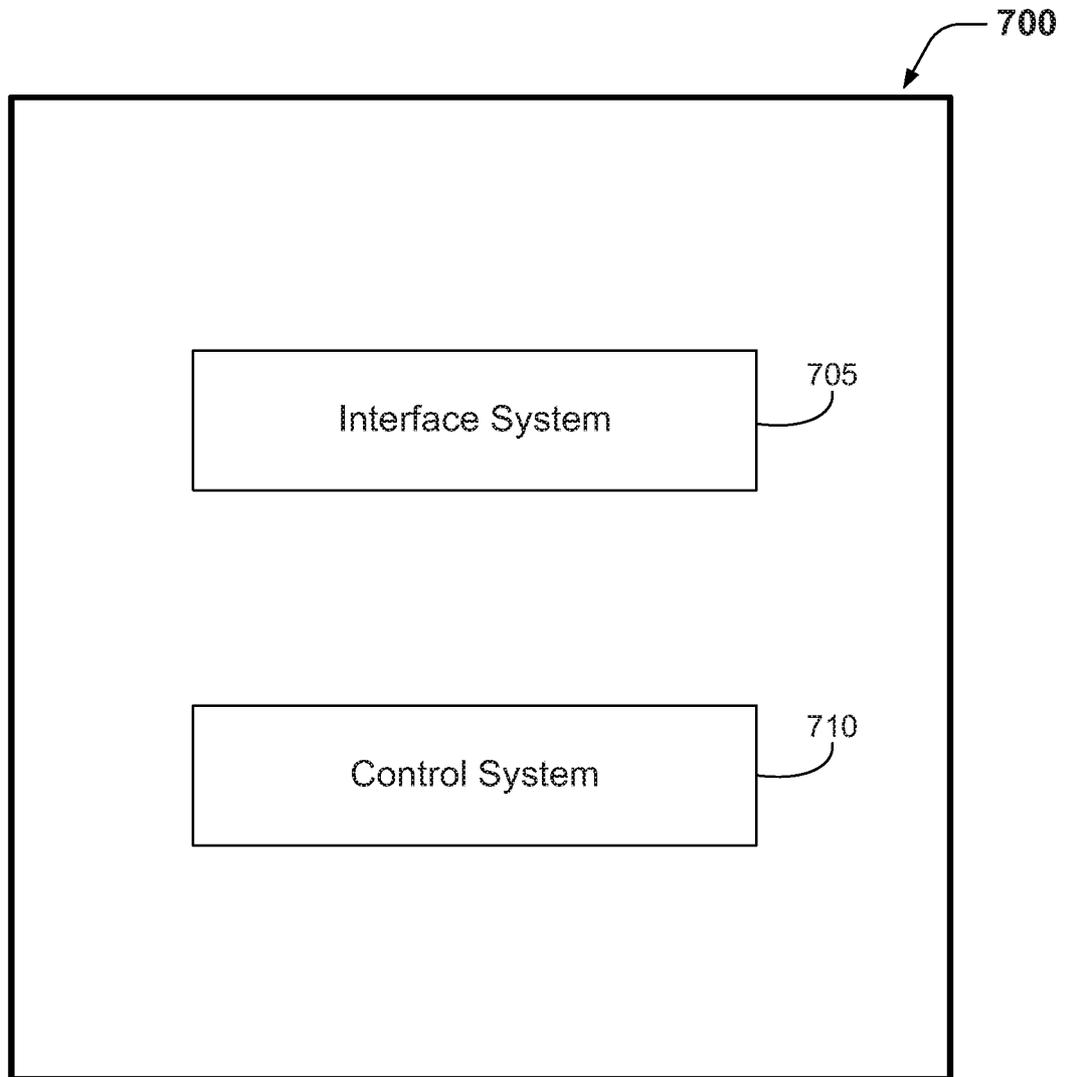
Figure 5E



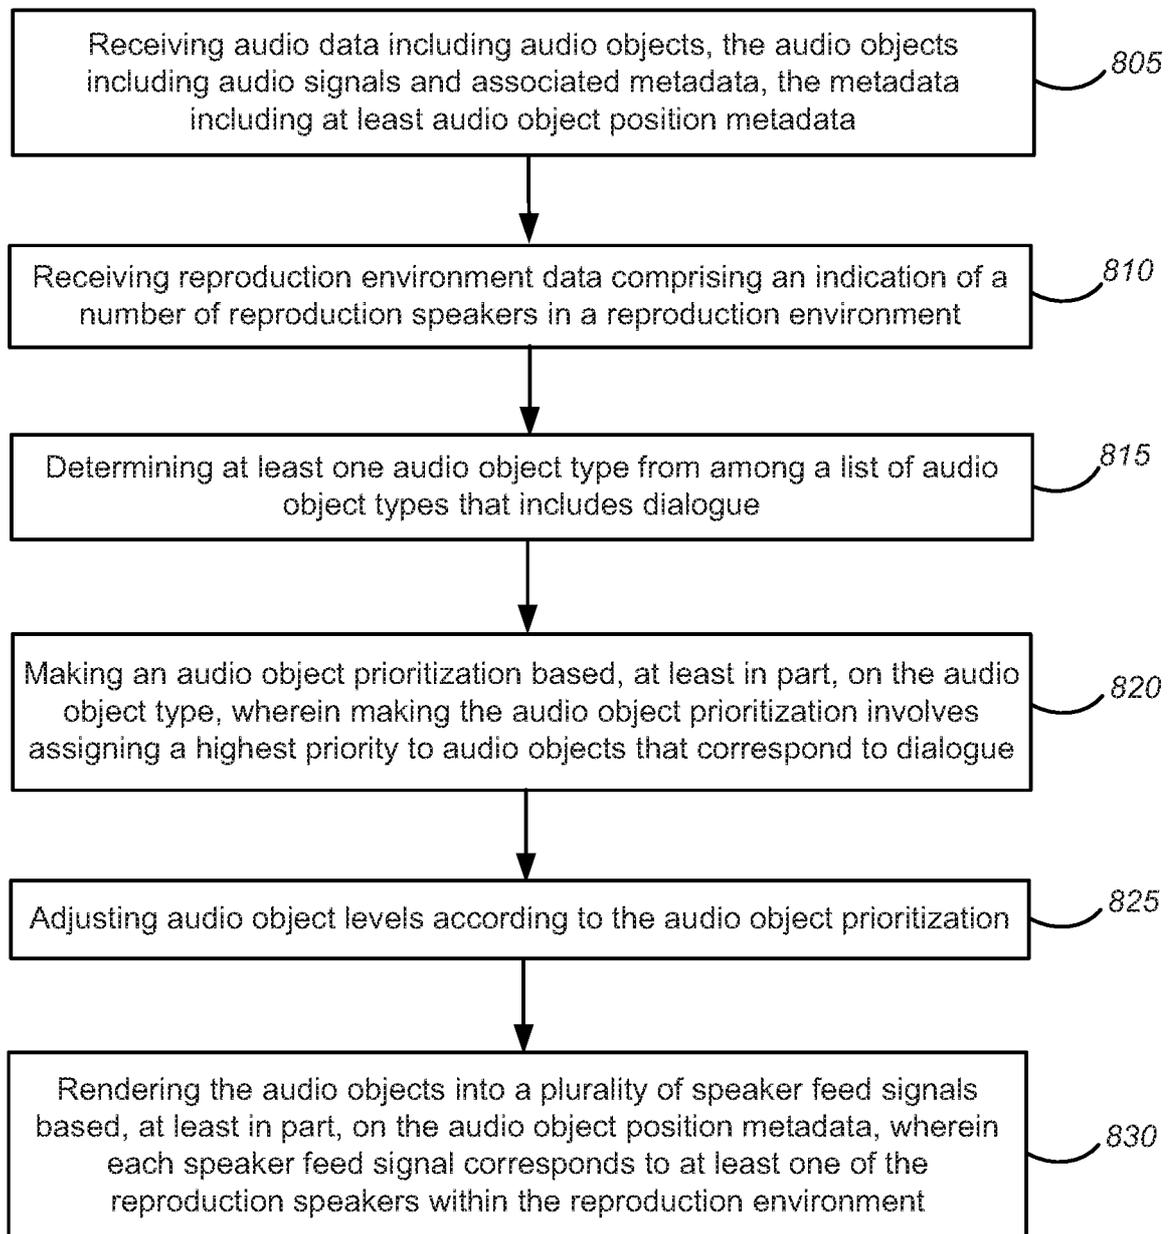
**Figure 6A**



**Figure 6B**

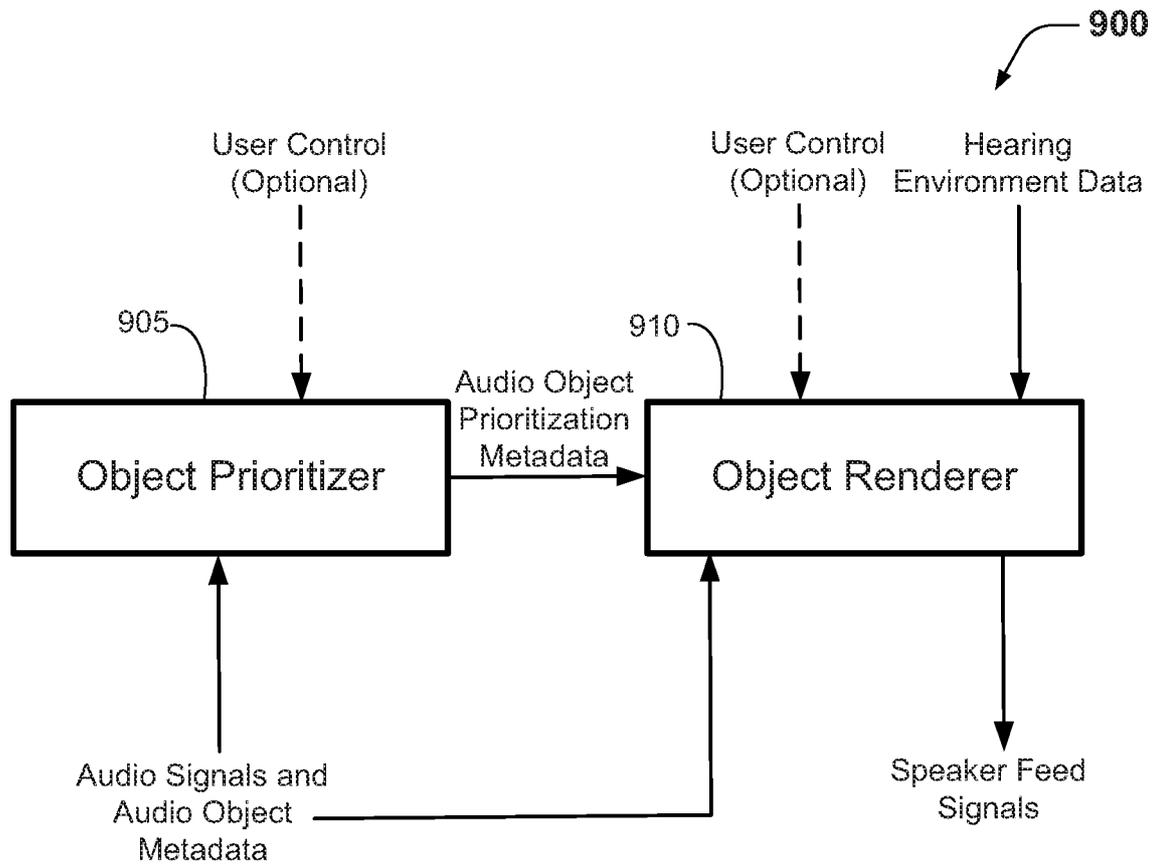


*Figure 7*



800 ↗

**Figure 8**



**Figure 9A**

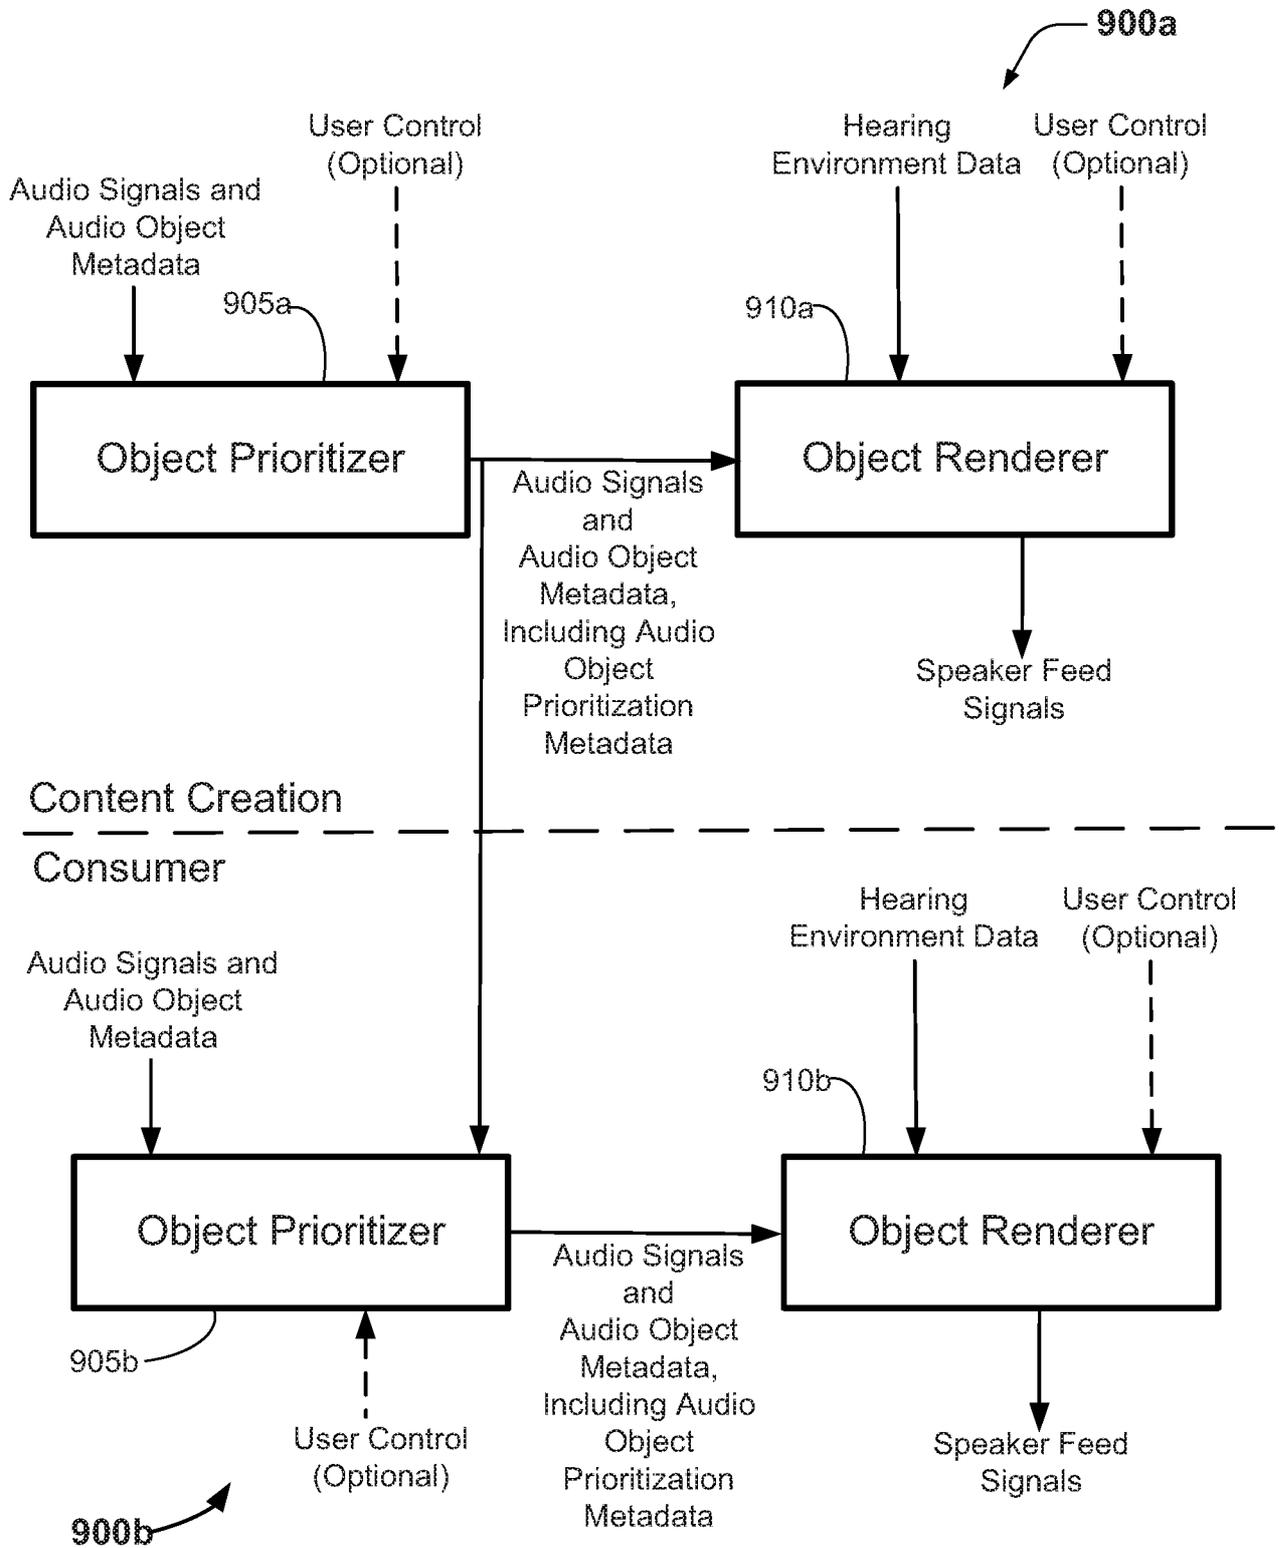
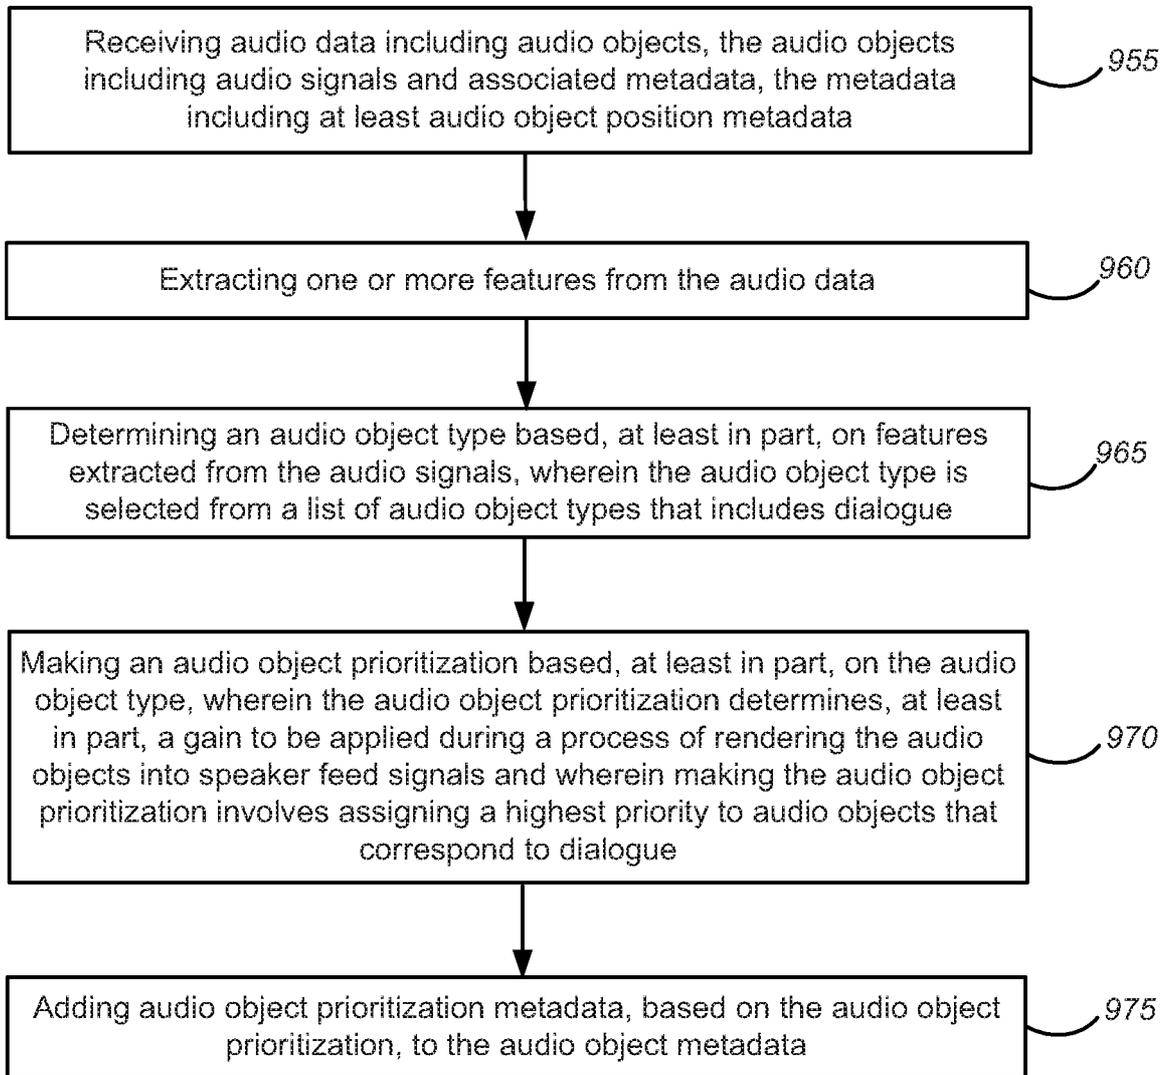


Figure 9B



950 ↗

**Figure 9C**

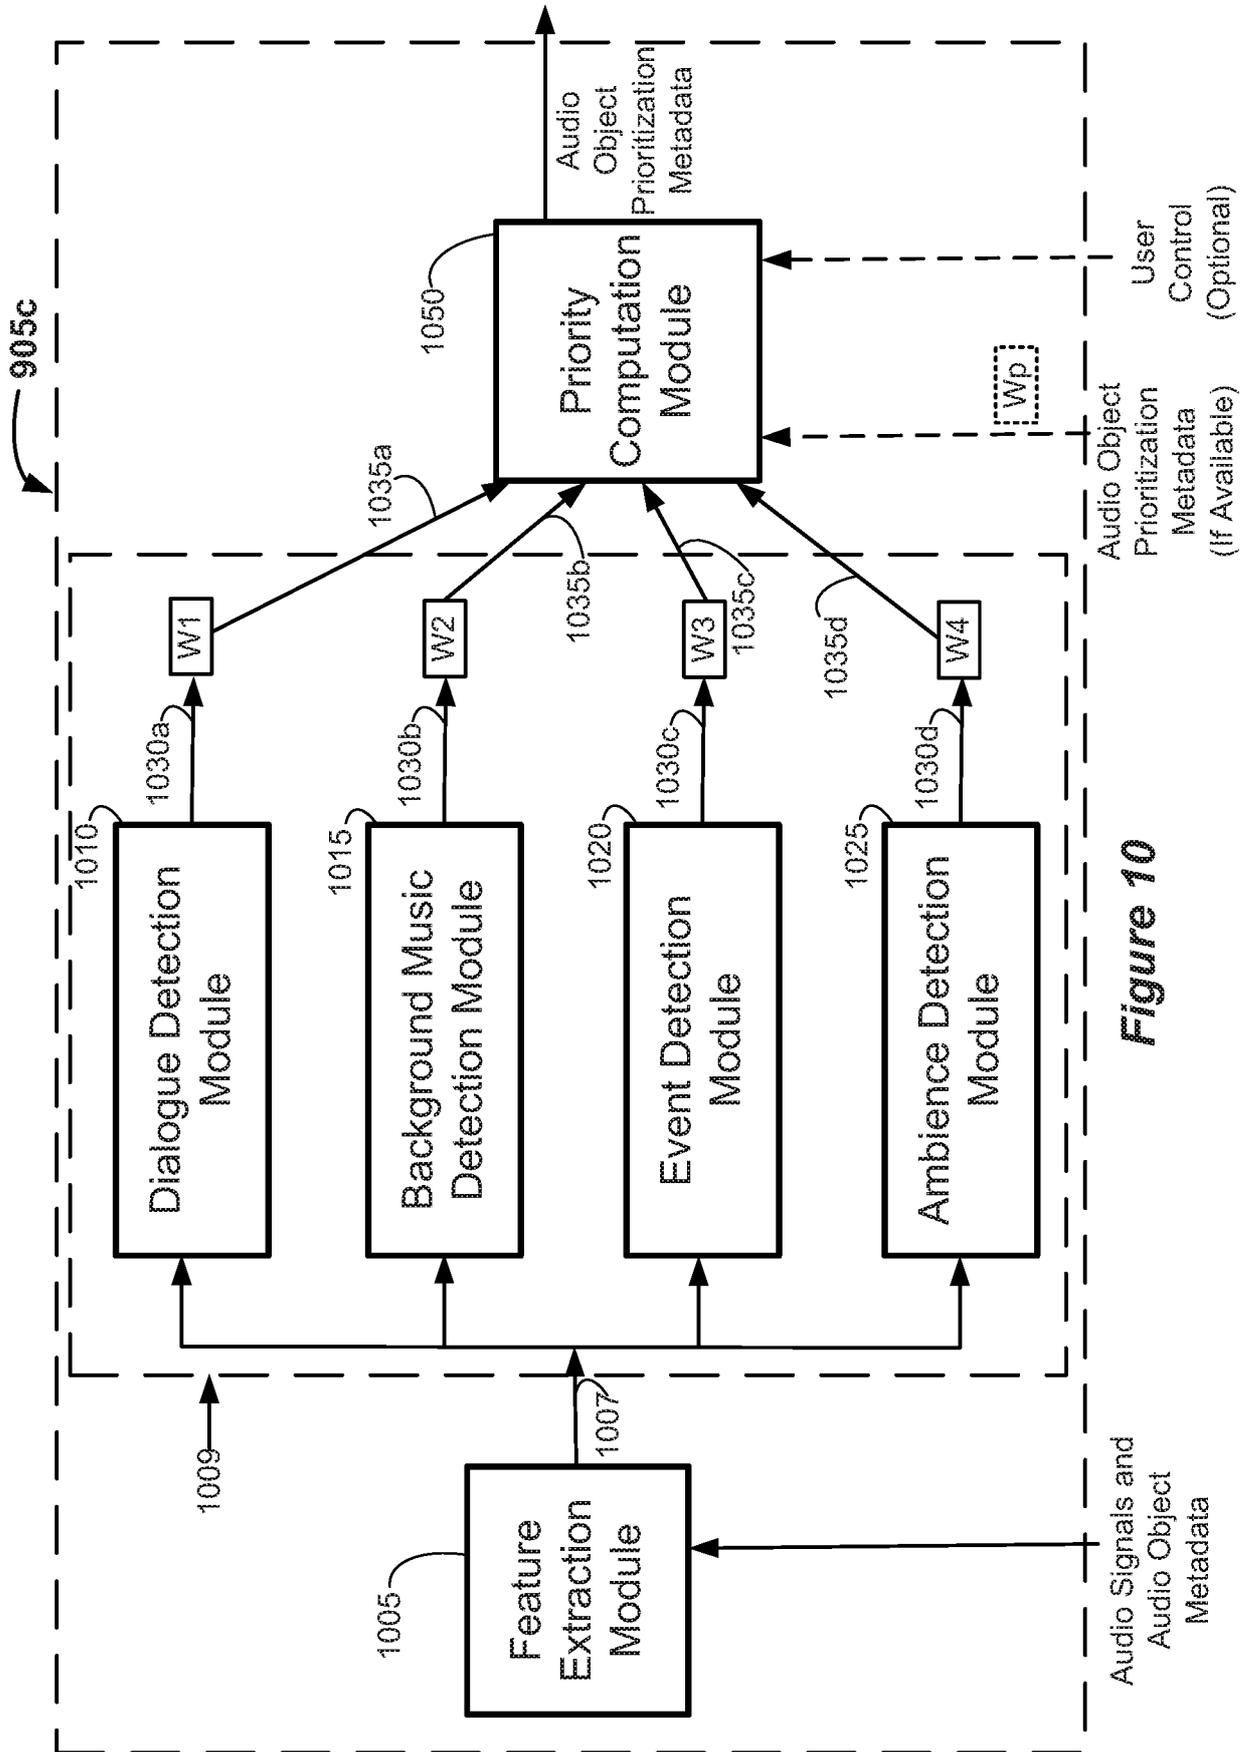
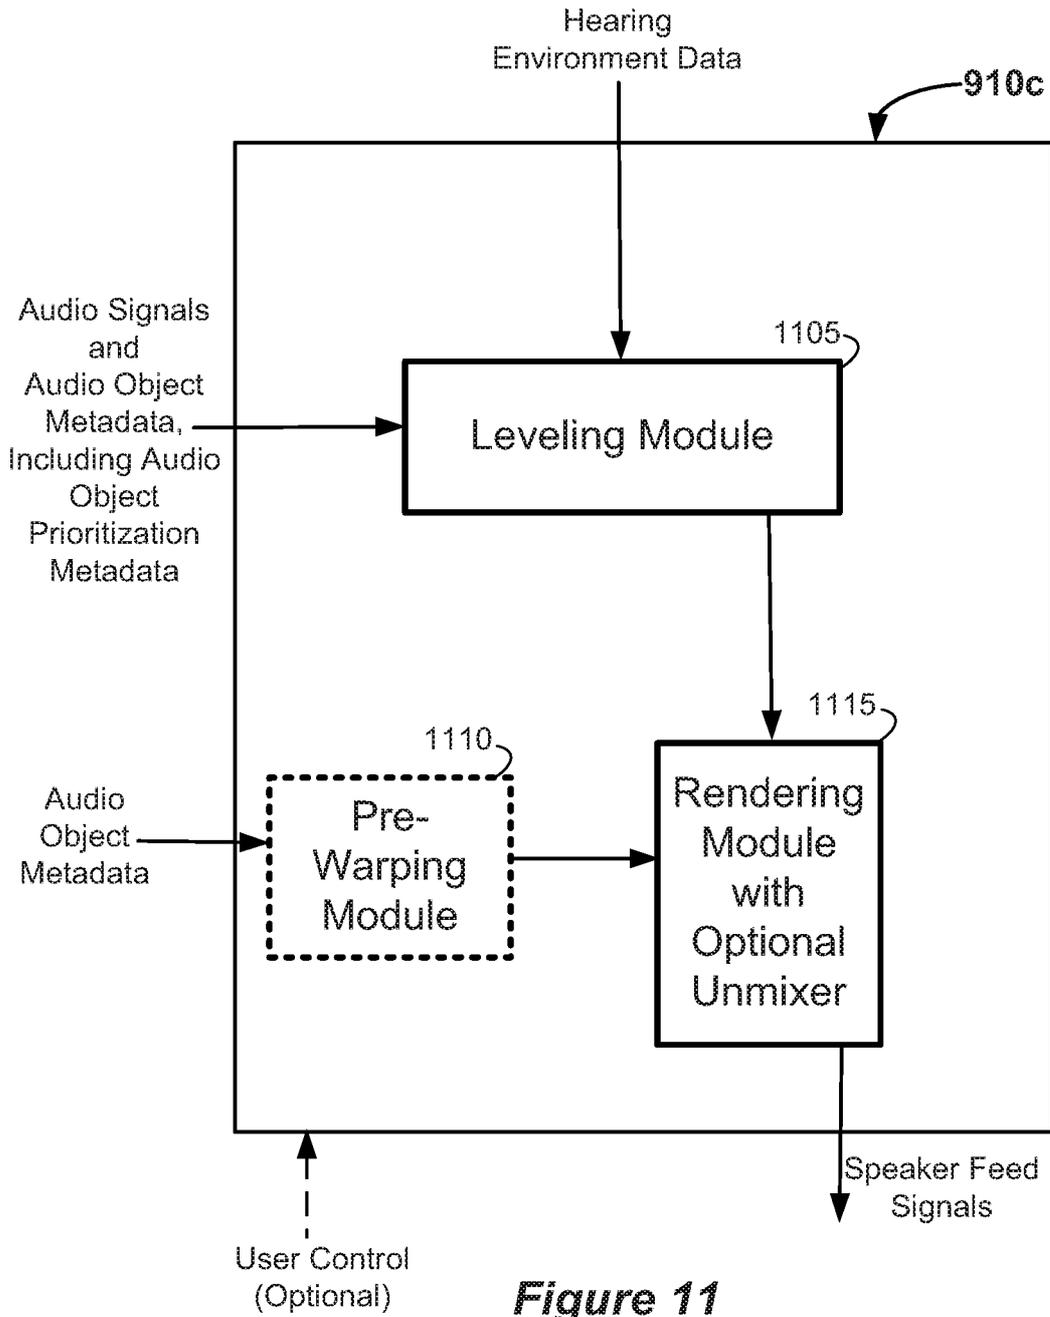


Figure 10



**Figure 11**

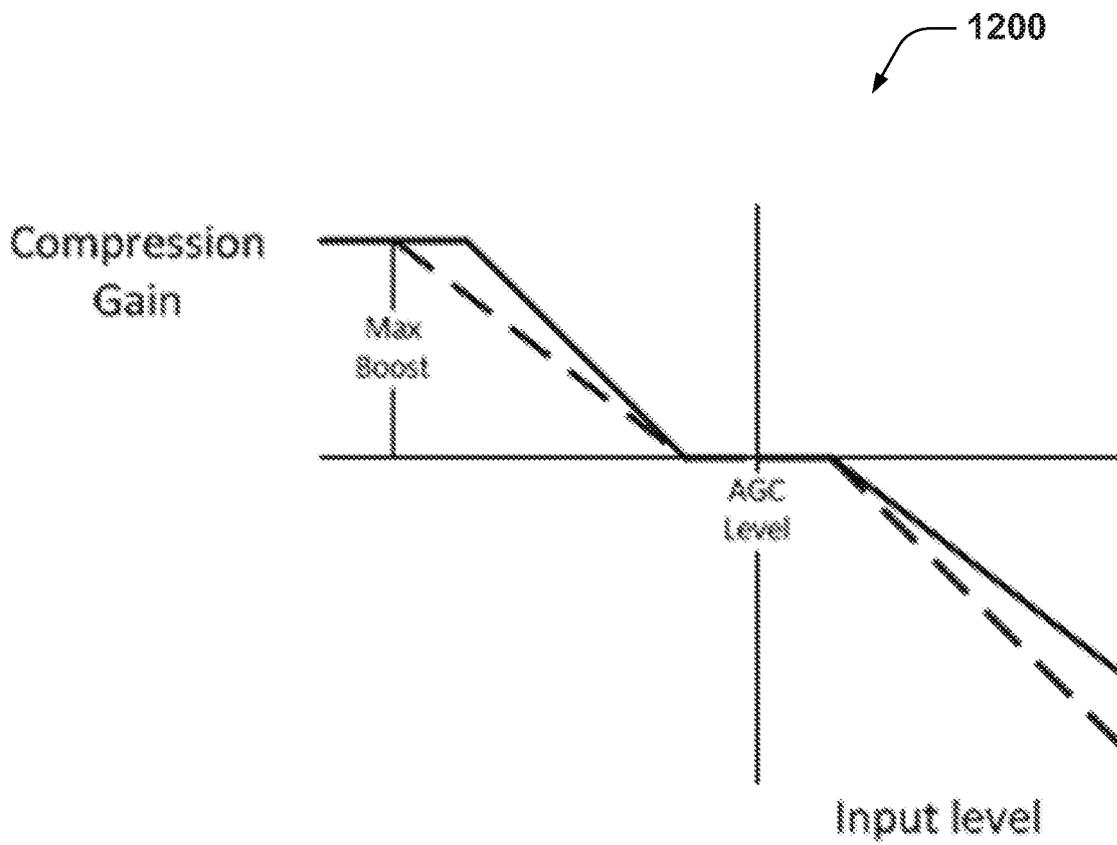


Figure 12

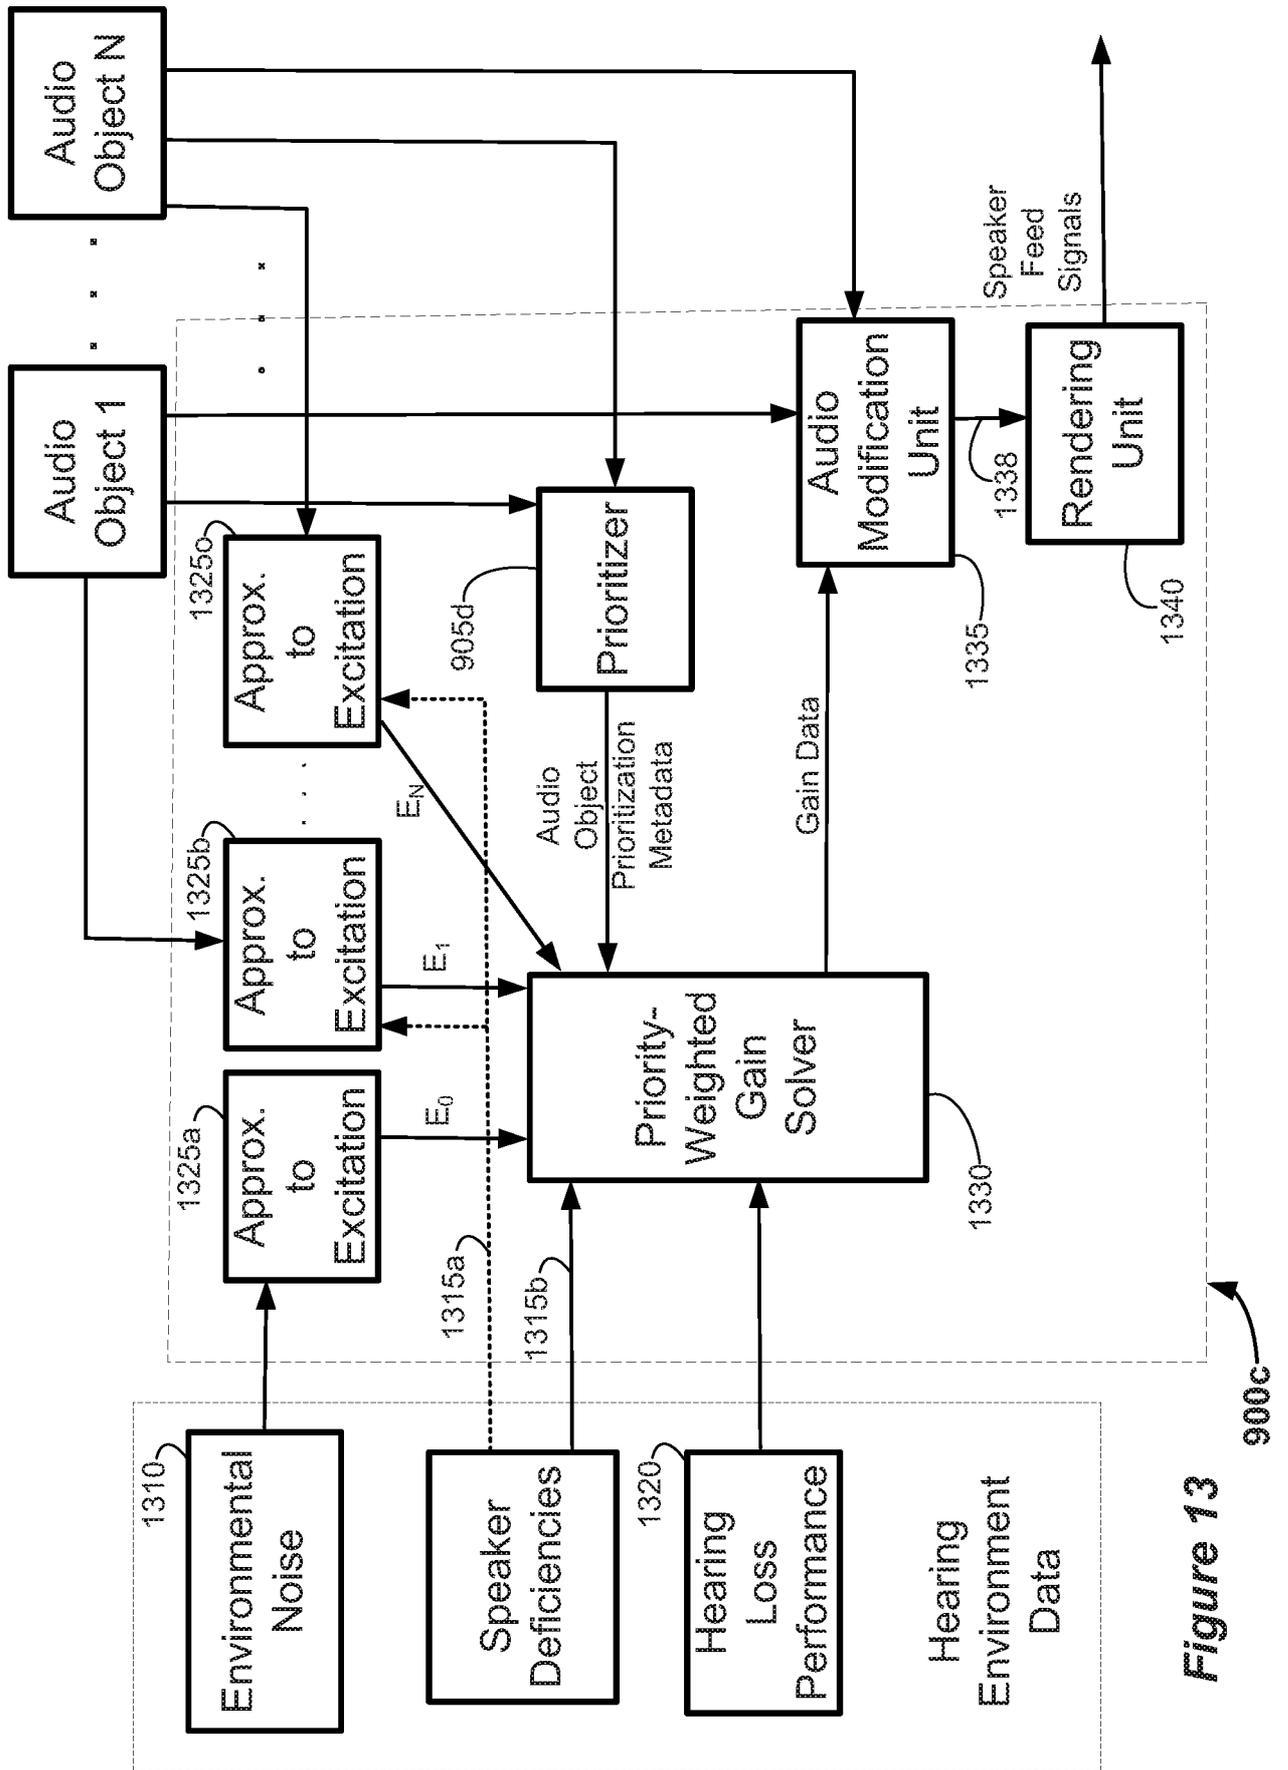


Figure 13

**INTERNATIONAL SEARCH REPORT**

International application No  
PCT/US2016/028295

<b>A. CLASSIFICATION OF SUBJECT MATTER</b> INV. H04S7/Q0 ADD. H04R27/O0                      H04R3/12		
According to International Patent Classification (IPC) or to both national classification and IPC		
<b>B. FIELDS SEARCHED</b>		
Minimum documentation searched (classification system followed by classification symbols) H04S H04R		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) EPO-Internal , INSPEC, WPI Data		
<b>C. DOCUMENTS CONSIDERED TO BE RELEVANT</b>		
<b>Category*</b>	<b>Citation of document, with indication, where appropriate, of the relevant passages</b>	<b>Relevant to claim No.</b>
Y	wo 2013/181272 A2 (DTS LLC [US] ) 5 December 2013 (2013-12-05) paragraphs [0025] , [0030] - [0034] , [0042] - [0044] -----	1-38
Y	us 7 974 422 B1 (HO CHI FAI [US] ET AL) 5 July 2011 (2011-07-05) col umn 4, l ines 7-34 col umn 5, l ines 21-47 col umn 7, l ines 21-57 ----- - / - -	1-38
<input checked="" type="checkbox"/> Further documents are listed in the continuation of Box C. <input checked="" type="checkbox"/> See patent family annex.		
* Special categories of cited documents :		
"A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier application or patent but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed		"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art "&" document member of the same patent family
Date of the actual completion of the international search  align="center">28 June 2016		Date of mailing of the international search report  align="center">07/07/2016
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016		Authorized officer  align="center">Rogała, Tomasz

**INTERNATIONAL SEARCH REPORT**

International application No PCT/US2016/028295
---

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	<p>Ken Pyle: "Do You Hear What I Hear? ", ' 19 June 2014 (2014-06-19) , XP055283553, Retrieved from the Internet: URL:http://video .com/2014/06/19/do-you-hear-what-i-hear/ [retrieved on 2016-06-24] page 1, lines 6-30</p> <p align="center">-----</p>	<p>2,22,23, 25,27 , 29,31 , 33-36,38</p>
A	<p>wo 2011/020065 AI (SRS LABS INC [US] ; TRACEY JAMES [US] ; KATSIANOS THEMIS [US] ; KRAEMER A) 17 February 2011 (2011-02-17) paragraphs [0077] - [0083]</p> <p align="center">-----</p>	<p>1-38</p>
A	<p>AU 2013 200 578 AI (FRAUNHOFER GES FORSCHUNG) 21 February 2013 (2013-02-21) abstract</p> <p align="center">-----</p>	<p>1-38</p>

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No PCT/US2016/028295
---

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
wo 2013181272 A2	05-12-2013	US 2013329922 AI Wo 2013181272 A2	12-12-2013 05-12-2013
-----			
US 7974422 BI	05-07-2011	US 7974422 BI US 2011285593 AI US 2012237005 AI US 2012250899 AI	05-07-2011 24-11-2011 20-09-2012 04-10-2012
-----			
wo 2011020065 AI	17-02-2011	CN 102549655 A CN 102576533 A EP 2465114 AI EP 2465259 AI JP 5635097 B2 JP 5726874 B2 JP 2013502183 A JP 2013502184 A KR 20120061869 A KR 20120062758 A US 2011040395 AI US 2011040396 AI US 2011040397 AI US 2013202129 AI wo 2011020065 AI wo 2011020067 AI	04--07--2012 11--07--2012 20--06--2012 20--06--2012 03-12-2014 03-06-2015 17--01--2013 17--01--2013 13--06-2012 14--06--2012 17--02--2011 17--02--2011 17--02--2011 08--08--2013 17--02--2011 17--02--2011
-----			
AU 2013200578 AI	21-02-2013	NONE	
-----			