



(51) International Patent Classification:

G06F 9/451 (2018.01) G06T 19/20 (2011.01)
G06F 3/04815 (2022.01) G06F 3/01 (2006.01)
G06F 3/04845 (2022.01) G06F 3/0481 (2022.01)
G06T 19/00 (2011.01) G06F 3/04842 (2022.01)

(21) International Application Number:

PCT/US2024/030462

(22) International Filing Date:

22 May 2024 (22.05.2024)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

63/471,257 05 June 2023 (05.06.2023) US
18/619,568 28 March 2024 (28.03.2024) US

(71) Applicant: APPLE INC. [US/US]; One Apple Park Way, Cupertino, California 95014 (US).

(72) Inventors: HANNEMANN, Raffael; c/o Apple Inc., One Apple Park Way, Cupertino, California 95014 (US). METROT, Sebastien; c/o Apple Inc., One Apple Park Way, Cupertino, California 95014 (US). MONSMA, Owen; c/o Apple Inc., One Apple Park Way, Cupertino, California 95014 (US). HAJAS, Peter L.; c/o Apple Inc., One

Apple Park Way, Cupertino, California 95014 (US). TURNER, James T.; c/o Apple Inc., One Apple Park Way, Cupertino, California 95014 (US).

(74) Agent: MORSE, Kyle; BakerHostetler, 600 Anton Blvd., Suite 900, COSTA MESA, California 92626 (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, CV, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IQ, IR, IS, IT, JM, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, MG, MK, MN, MU, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, CV, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SC, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, ME, MK, MT, NL, NO, PL, PT, RO, RS, SE,

(54) Title: TECHNIQUES FOR MANAGING COMPUTER-GENERATED EXPERIENCES

(57) Abstract: Some techniques are described herein for managing computer-generated environments, including methods for managing the size of virtual objects and managing an experience.




FIG. 4

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

Published:

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the
claims and to be republished in the event of receipt of
amendments (Rule 48.2(h))*

TECHNIQUES FOR MANAGING COMPUTER-GENERATED EXPERIENCES**CROSS-REFERENCE TO RELATED APPLICATIONS**

[0001] The present application claims priority to U.S. Non-Provisional Patent Application Serial No. 18/619,568, entitled “TECHNIQUES FOR MANAGING COMPUTER-GENERATED EXPERIENCES” filed March 28, 2024, and to U.S. Provisional Patent Application Serial No. 63/471,257 entitled “TECHNIQUES FOR MANAGING COMPUTER-GENERATED EXPERIENCES,” filed June 5, 2023, which are hereby incorporated by reference in their entireties for all purposes.

BACKGROUND

[0002] Today, people are using electronic device to attend more computer-generated experiences, including working, playing games, or attending events. There is a need to manage the computer-generated experiences to provide more value to more people.

SUMMARY

[0003] Current techniques for managing computer-generated experiences are generally ineffective and/or inefficient. For example, some techniques require software developers to program solutions to manage computer-generated experiences independently, which can be time consuming and ineffective. This disclosure provides more effective and/or efficient techniques for managing computer-generated experiences for a computer system that can display virtual objects in immersive virtual environments. It should be recognized that other types of electronic devices can be used with techniques described herein. For example, computer systems that cannot display virtual objects in immersive virtual environments can be managed using the techniques described herein. In addition, techniques optionally complement or replace other techniques for managing three-dimensional environments.

[0004] Some techniques are described herein for managing the size of virtual objects as an environment or user interface that includes the virtual objects is resized. Such techniques provide techniques for maintaining the physical size of the virtual object or maintaining the actual size of the virtual object as the environment or user interface is resized.

[0005] In some examples, a method that is performed by a computer system is described. In some examples, the method comprises: receiving, from a first application, a request to render a first object in an environment; and in response to receiving the request to render the first object in the environment, rendering the first object; after rendering the first object and while the first object is being displayed in the environment, detecting an indication of a request to change a size of content provided by the first application; and in response to detecting the indication of the request to change the size of content provided by the first application, rendering a representation of the environment by: in accordance with a determination that the environment is rendered according to a first mode and that the first object is a first type of object, wherein the environment corresponds to the first application and a second application different from the first application while the environment is rendered according to the first mode, maintaining a physical size of the first object without maintaining an angular size of the first object; in accordance with a determination that the environment is rendered according to the first mode and that the first object is a second type of object different from the first type of object, maintaining the angular size of the first object without maintaining the physical size of the first object; and in accordance with a determination that the environment is rendered according to a second mode, wherein the environment corresponds to the first application without corresponding to another application different from the first application while the environment is rendered according to the second mode, maintaining the physical size of the first object without maintaining the angular size of the first object.

[0006] In some examples, a non-transitory computer-readable storage medium storing one or more programs configured to be executed by one or more processors of a computer system is described. In some examples, the one or more programs includes instructions for: receiving, from a first application, a request to render a first object in an environment; and in response to receiving the request to render the first object in the environment, rendering the first object; after rendering the first object and while the first object is being displayed in the environment, detecting an indication of a request to change a size of content provided by the first application; and in response to detecting the indication of the request to change the size of content provided by the first application, rendering a representation of the environment by: in accordance with a determination that the environment is rendered according to a first mode and that the first object is a first type of object, wherein the environment corresponds to the first application and a second application different from the first application while the

environment is rendered according to the first mode, maintaining a physical size of the first object without maintaining an angular size of the first object; in accordance with a determination that the environment is rendered according to the first mode and that the first object is a second type of object different from the first type of object, maintaining the angular size of the first object without maintaining the physical size of the first object; and in accordance with a determination that the environment is rendered according to a second mode, wherein the environment corresponds to the first application without corresponding to another application different from the first application while the environment is rendered according to the second mode, maintaining the physical size of the first object without maintaining the angular size of the first object.

[0007] In some examples, a transitory computer-readable storage medium storing one or more programs configured to be executed by one or more processors of a computer system is described. In some examples, the one or more programs includes instructions for: receiving, from a first application, a request to render a first object in an environment; and in response to receiving the request to render the first object in the environment, rendering the first object; after rendering the first object and while the first object is being displayed in the environment, detecting an indication of a request to change a size of content provided by the first application; and in response to detecting the indication of the request to change the size of content provided by the first application, rendering a representation of the environment by: in accordance with a determination that the environment is rendered according to a first mode and that the first object is a first type of object, wherein the environment corresponds to the first application and a second application different from the first application while the environment is rendered according to the first mode, maintaining a physical size of the first object without maintaining an angular size of the first object; in accordance with a determination that the environment is rendered according to the first mode and that the first object is a second type of object different from the first type of object, maintaining the angular size of the first object without maintaining the physical size of the first object; and in accordance with a determination that the environment is rendered according to a second mode, wherein the environment corresponds to the first application without corresponding to another application different from the first application while the environment is rendered according to the second mode, maintaining the physical size of the first object without maintaining the angular size of the first object.

[0008] In some examples, a computer system comprising one or more processors and memory storing one or more programs configured to be executed by the one or more processors is described. In some examples, the one or more programs includes instructions for: receiving, from a first application, a request to render a first object in an environment; and in response to receiving the request to render the first object in the environment, rendering the first object; after rendering the first object and while the first object is being displayed in the environment, detecting an indication of a request to change a size of content provided by the first application; and in response to detecting the indication of the request to change the size of content provided by the first application, rendering a representation of the environment by: in accordance with a determination that the environment is rendered according to a first mode and that the first object is a first type of object, wherein the environment corresponds to the first application and a second application different from the first application while the environment is rendered according to the first mode, maintaining a physical size of the first object without maintaining an angular size of the first object; in accordance with a determination that the environment is rendered according to the first mode and that the first object is a second type of object different from the first type of object, maintaining the angular size of the first object without maintaining the physical size of the first object; and in accordance with a determination that the environment is rendered according to a second mode, wherein the environment corresponds to the first application without corresponding to another application different from the first application while the environment is rendered according to the second mode, maintaining the physical size of the first object without maintaining the angular size of the first object.

[0009] In some examples, a computer system is comprising means for performing each of the following steps: receiving, from a first application, a request to render a first object in an environment; and in response to receiving the request to render the first object in the environment, rendering the first object; after rendering the first object and while the first object is being displayed in the environment, detecting an indication of a request to change a size of content provided by the first application; and in response to detecting the indication of the request to change the size of content provided by the first application, rendering a representation of the environment by: in accordance with a determination that the environment is rendered according to a first mode and that the first object is a first type of object, wherein the environment corresponds to the first application and a second application different from the first application while the environment is rendered according to the first

mode, maintaining a physical size of the first object without maintaining an angular size of the first object; in accordance with a determination that the environment is rendered according to the first mode and that the first object is a second type of object different from the first type of object, maintaining the angular size of the first object without maintaining the physical size of the first object; and in accordance with a determination that the environment is rendered according to a second mode, wherein the environment corresponds to the first application without corresponding to another application different from the first application while the environment is rendered according to the second mode, maintaining the physical size of the first object without maintaining the angular size of the first object.

[0010] In some examples, a computer program product is described. In some examples, the computer program product comprises one or more programs configured to be executed by one or more processors of a computer system. In some examples, the one or more programs include instructions for: receiving, from a first application, a request to render a first object in an environment; and in response to receiving the request to render the first object in the environment, rendering the first object; after rendering the first object and while the first object is being displayed in the environment, detecting an indication of a request to change a size of content provided by the first application; and in response to detecting the indication of the request to change the size of content provided by the first application, rendering a representation of the environment by: in accordance with a determination that the environment is rendered according to a first mode and that the first object is a first type of object, wherein the environment corresponds to the first application and a second application different from the first application while the environment is rendered according to the first mode, maintaining a physical size of the first object without maintaining an angular size of the first object; in accordance with a determination that the environment is rendered according to the first mode and that the first object is a second type of object different from the first type of object, maintaining the angular size of the first object without maintaining the physical size of the first object; and in accordance with a determination that the environment is rendered according to a second mode, wherein the environment corresponds to the first application without corresponding to another application different from the first application while the environment is rendered according to the second mode, maintaining the physical size of the first object without maintaining the angular size of the first object.

[0011] In some examples, a method that is performed by a computer system is described. In some examples, the method comprises: sending, from an application, a request for a pose associated with a user with respect to an origin of an experience for the application; and after sending the request for the pose associated with the user with respect to the origin of the experience for the application: in accordance with a determination that a first number of users are in the experience for the application: receiving a first pose; and rendering the experience for the application based on the first pose; and in accordance with a determination that a second number of users are in the experience for the application: receiving a second pose different from the first pose; and rendering the experience for the application based on the second pose.

[0012] In some examples, a non-transitory computer-readable storage medium storing one or more programs configured to be executed by one or more processors of a computer system is described. In some examples, the one or more programs includes instructions for: sending, from an application, a request for a pose associated with a user with respect to an origin of an experience for the application; and after sending the request for the pose associated with the user with respect to the origin of the experience for the application: in accordance with a determination that a first number of users are in the experience for the application: receiving a first pose; and rendering the experience for the application based on the first pose; and in accordance with a determination that a second number of users are in the experience for the application: receiving a second pose different from the first pose; and rendering the experience for the application based on the second pose.

[0013] In some examples, a transitory computer-readable storage medium storing one or more programs configured to be executed by one or more processors of a computer system is described. In some examples, the one or more programs includes instructions for: sending, from an application, a request for a pose associated with a user with respect to an origin of an experience for the application; and after sending the request for the pose associated with the user with respect to the origin of the experience for the application: in accordance with a determination that a first number of users are in the experience for the application: receiving a first pose; and rendering the experience for the application based on the first pose; and in accordance with a determination that a second number of users are in the experience for the application: receiving a second pose different from the first pose; and rendering the experience for the application based on the second pose.

[0014] In some examples, a computer system comprising one or more processors and memory storing one or more programs configured to be executed by the one or more processors is described. In some examples, the one or more programs includes instructions for: sending, from an application, a request for a pose associated with a user with respect to an origin of an experience for the application; and after sending the request for the pose associated with the user with respect to the origin of the experience for the application: in accordance with a determination that a first number of users are in the experience for the application: receiving a first pose; and rendering the experience for the application based on the first pose; and in accordance with a determination that a second number of users are in the experience for the application: receiving a second pose different from the first pose; and rendering the experience for the application based on the second pose.

[0015] In some examples, a computer system is comprising means for performing each of the following steps: sending, from an application, a request for a pose associated with a user with respect to an origin of an experience for the application; and after sending the request for the pose associated with the user with respect to the origin of the experience for the application: in accordance with a determination that a first number of users are in the experience for the application: receiving a first pose; and rendering the experience for the application based on the first pose; and in accordance with a determination that a second number of users are in the experience for the application: receiving a second pose different from the first pose; and rendering the experience for the application based on the second pose.

[0016] In some examples, a computer program product is described. In some examples, the computer program product comprises one or more programs configured to be executed by one or more processors of a computer system. In some examples, the one or more programs include instructions for: sending, from an application, a request for a pose associated with a user with respect to an origin of an experience for the application; and after sending the request for the pose associated with the user with respect to the origin of the experience for the application: in accordance with a determination that a first number of users are in the experience for the application: receiving a first pose; and rendering the experience for the application based on the first pose; and in accordance with a determination that a second number of users are in the experience for the application: receiving a second pose different

from the first pose; and rendering the experience for the application based on the second pose.

[0017] Executable instructions for performing these functions are, optionally, included in a non-transitory computer-readable storage medium or other computer program product configured for execution by one or more processors. Executable instructions for performing these functions are, optionally, included in a transitory computer-readable storage medium or other computer program product configured for execution by one or more processors.

DESCRIPTION OF THE FIGURES

[0018] For a better understanding of the various described examples, reference should be made to the Detailed Description below, in conjunction with the following drawings in which like reference numerals refer to corresponding parts throughout the figures.

[0019] FIG. 1 illustrates an example system architecture including various electronic devices that may implement the subject system in accordance with some examples.

[0020] FIG. 2 illustrates a block diagram of example features of an electronic device in accordance with some examples.

[0021] FIG. 3 is a block diagram illustrating a computer system in accordance with some examples.

[0022] FIG. 4 is a block diagram for managing the size of virtual objects in accordance with some examples.

[0023] FIG. 5 is a flow diagram illustrating a method for managing the size of virtual objects in accordance with some examples.

[0024] FIGS. 6A-6B are user interfaces for managing an experience in accordance with some examples.

[0025] FIG. 7 is a flow diagram illustrating a method for managing an experience in accordance with some examples.

[0026] FIG. 8 illustrates an electronic system with which some examples of the subject technology may be implemented in accordance with some examples.

DETAILED DESCRIPTION

[0027] The detailed description set forth below is intended as a description of various configurations of the subject technology and is not intended to represent the only configurations in which the subject technology can be practiced. The appended drawings are incorporated herein and constitute a part of the detailed description. The detailed description includes specific details for the purpose of providing a thorough understanding of the subject technology. However, the subject technology is not limited to the specific details set forth herein and can be practiced using one or more other examples. In some examples, structures and components are shown in block diagram form in order to avoid obscuring the concepts of the subject technology.

[0028] Methods and/or processes described herein can include one or more steps that are contingent upon one or more conditions being satisfied. It should be understood that a method can occur over multiple iterations of the same process with different steps of the method being satisfied in different iterations. For example, if a method requires performing a first step upon a determination that a set of one or more criteria is met and a second step upon a determination that the set of one or more criteria is not met, a person of ordinary skill in the art would appreciate that the steps of the method are repeated until both conditions, in no particular order, are satisfied. Thus, a method described with steps that are contingent upon a condition being satisfied can be rewritten as a method that is repeated until each of the conditions described in the method are satisfied. This, however, is not required of system or computer readable medium claims where the system or computer readable medium claims include instructions for performing one or more steps that are contingent upon one or more conditions being satisfied. Because the instructions for the system or computer readable medium claims are stored in one or more processors and/or at one or more memory locations, the system or computer readable medium claims include logic that can determine whether the one or more conditions have been satisfied without explicitly repeating steps of a method until all of the conditions upon which steps in the method are contingent have been satisfied. A person having ordinary skill in the art would also understand that, similar to a method with contingent steps, a system or computer readable storage medium can repeat the steps of a method as many times as needed to ensure that all of the contingent steps have been performed.

[0029] Although the following description uses terms “first,” “second,” “third,” etc. to describe various elements, these elements should not be limited by the terms. In some examples, these terms are used to distinguish one element from another. For example, a first subsystem could be termed a second subsystem, and, similarly, a subsystem device could be termed a subsystem device, without departing from the scope of the various described examples. In some examples, the first subsystem and the second subsystem are two separate references to the same subsystem. In some examples, the first subsystem and the second subsystem are both subsystems, but they are not the same subsystem or the same type of subsystem.

[0030] The terminology used in the description of the various described examples herein is for the purpose of describing particular examples only and is not intended to be limiting. As used in the description of the various described examples and the appended claims, the singular forms “a,” “an,” and “the” are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will also be understood that the term “and/or” as used herein refers to and encompasses any and all possible combinations of one or more of the associated listed items. It will be further understood that the terms “includes,” “including,” “comprises,” and/or “comprising,” when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof.

[0031] The term “if” is, optionally, construed to mean “when,” “upon,” “in response to determining,” “in response to detecting,” or “in accordance with a determination that” depending on the context. Similarly, the phrase “if it is determined” or “if [a stated condition or event] is detected” is, optionally, construed to mean “upon determining,” “in response to determining,” “upon detecting [the stated condition or event],” “in response to detecting [the stated condition or event],” or “in accordance with a determination that [the stated condition or event]” depending on the context.

[0032] A physical environment refers to a physical world that people can sense and/or interact with without aid of electronic devices. The physical environment may include physical features such as a physical surface or a physical object. For example, the physical environment corresponds to a physical park that includes physical trees, physical buildings, and physical people. People can directly sense and/or interact with the physical environment

such as through sight, touch, hearing, taste, and smell. In contrast, an extended reality (XR) environment refers to a wholly or partially simulated environment that people sense and/or interact with via an electronic device. For example, the XR environment may include augmented reality (AR) content, mixed reality (MR) content, virtual reality (VR) content, and/or the like. With an XR system, a subset of a person's physical motions, or representations thereof, are tracked, and, in response, one or more characteristics of one or more virtual objects simulated in the XR environment are adjusted in a manner that comports with at least one law of physics. As one example, the XR system may detect head movement and, in response, adjust graphical content and an acoustic field presented to the person in a manner similar to how such views and sounds would change in a physical environment. As another example, the XR system may detect movement of the electronic device presenting the XR environment (e.g., a mobile phone, a tablet, a laptop, or the like) and, in response, adjust graphical content and an acoustic field presented to the person in a manner similar to how such views and sounds would change in a physical environment. In some situations (e.g., for accessibility reasons), the XR system may adjust characteristic(s) of graphical content in the XR environment in response to representations of physical motions (e.g., vocal commands).

[0033] There are many different types of electronic systems that enable a person to sense and/or interact with various XR environments. Examples include head mountable systems, projection-based systems, heads-up displays (HUDs), vehicle windshields having integrated display capability, windows having integrated display capability, displays formed as lenses designed to be placed on a person's eyes (e.g., similar to contact lenses), headphones/earphones, speaker arrays, input systems (e.g., wearable or handheld controllers with or without haptic feedback), smartphones, tablets, and desktop/laptop computers. A head mountable system may have one or more speaker(s) and an integrated opaque display. Alternatively, a head mountable system may be configured to accept an external opaque display (e.g., a smartphone). The head mountable system may incorporate one or more imaging sensors to capture images or video of the physical environment, and/or one or more microphones to capture audio of the physical environment. Rather than an opaque display, a head mountable system may have a transparent or translucent display. The transparent or translucent display may have a medium through which light representative of images is directed to a person's eyes. The display may utilize digital light projection, OLEDs, LEDs, uLEDs, liquid crystal on silicon, laser scanning light source, or any combination of these technologies. The medium may be an optical waveguide, a hologram medium, an optical

combiner, an optical reflector, or any combination thereof. In some examples, the transparent or translucent display may be configured to become opaque selectively. Projection-based systems may employ retinal projection technology that projects graphical images onto a person's retina. Projection systems also may be configured to project virtual objects into the physical environment, for example, as a hologram or on a physical surface.

[0034] FIG. 1 illustrates an example system architecture 100 including various electronic devices that may implement the subject system in accordance with some examples. Not all of the depicted components may be used in all examples, however, and some examples may include additional or different components than those shown in the figure. Variations in the arrangement and type of the components may be made without departing from the spirit or scope of the claims as set forth herein. Additional components, different components, or fewer components may be provided.

[0035] The system architecture 100 includes an electronic device 105, a handheld electronic device 104, an electronic device 110, an electronic device 115, and a server 120. For explanatory purposes, the system architecture 100 is illustrated in FIG. 1 as including the electronic device 105, the handheld electronic device 104, the electronic device 110, the electronic device 115, and the server 120; however, the system architecture 100 may include any number of electronic devices, and any number of servers or a data center including multiple servers.

[0036] The electronic device 105 may be implemented, for example, as a tablet device, a smartphone, or as a head mountable portable system (e.g., worn by a user 101). The electronic device 105 includes a display system capable of presenting a visualization of an extended reality environment to the user. The electronic device 105 may be powered with a battery and/or another power supply. In an example, the display system of the electronic device 105 provides a stereoscopic presentation of the extended reality environment, enabling a three-dimensional visual display of a rendering of a particular scene, to the user. In some examples, instead of, or in addition to, utilizing the electronic device 105 to access an extended reality environment, the user may use a handheld electronic device 104, such as a tablet, watch, mobile device, and the like.

[0037] The electronic device 105 may include one or more cameras such as camera(s) 150 (e.g., visible light cameras, infrared cameras, etc.) For example, the electronic device 105

may include multiple cameras 150. For example, the multiple cameras 150 may include a left facing camera, a front facing camera, a right facing camera, a down facing camera, a left-down facing camera, a right-down facing camera, an up facing camera, one or more eye-facing cameras, and/or other cameras. Each of the cameras 150 may include one or more image sensors (e.g., charged coupled device (CCD) image sensors, complementary metal oxide semiconductor (CMOS) image sensors, or the like).

[0038] Further, the electronic device 105 may include various sensors 152 including, but not limited to, other cameras, other image sensors, touch sensors, microphones, inertial measurement units (IMU), heart rate sensors, temperature sensors, depth sensors (e.g., Lidar sensors, radar sensors, sonar sensors, time-of-flight sensors, etc.), GPS sensors, Wi-Fi sensors, near-field communications sensors, radio frequency sensors, etc. Moreover, the electronic device 105 may include hardware elements that can receive user input such as hardware buttons or switches. User inputs detected by such cameras, sensors, and/or hardware elements may correspond to, for example, various input modalities. For example, such input modalities may include, but are not limited to, facial tracking, eye tracking (e.g., gaze direction), hand tracking, gesture tracking, biometric readings (e.g., heart rate, pulse, pupil dilation, breath, temperature, electroencephalogram, olfactory), recognizing speech or audio (e.g., particular hotwords), and activating buttons or switches, etc. In some examples, facial tracking, gaze tracking, hand tracking, gesture tracking, object tracking, and/or physical environment mapping processes (e.g., system processes and/or application processes) may utilize images (e.g., image frames) captured by one or more image sensors of the cameras 150 and/or the sensors 152.

[0039] In some examples, the electronic device 105 may be communicatively coupled to a base device such as the electronic device 110 and/or the electronic device 115. Such a base device may, in general, include more computing resources and/or available power in comparison with the electronic device 105. In an example, the electronic device 105 may operate in various modes. For instance, the electronic device 105 can operate in a standalone mode independent of any base device. When the electronic device 105 operates in the standalone mode, the number of input modalities may be constrained by power and/or processing limitations of the electronic device 105 such as available battery power of the device. In response to power limitations, the electronic device 105 may deactivate certain sensors within the device itself to preserve battery power and/or to free processing resources.

[0040] The electronic device 105 may also operate in a wireless tethered mode (e.g., connected via a wireless connection with a base device), working in conjunction with a given base device. The electronic device 105 may also work in a connected mode where the electronic device 105 is physically connected to a base device (e.g., via a cable or some other physical connector) and may utilize power resources provided by the base device (e.g., where the base device is charging the electronic device 105 and/or providing power to the electronic device 105 while physically connected).

[0041] When the electronic device 105 operates in the wireless tethered mode or the connected mode, a least a portion of processing user inputs and/or rendering the extended reality environment may be offloaded to the base device thereby reducing processing burdens on the electronic device 105. For instance, in an example, the electronic device 105 works in conjunction with the electronic device 110 or the electronic device 115 to generate an extended reality environment including physical and/or virtual objects that enables different forms of interaction (e.g., visual, auditory, and/or physical or tactile interaction) between the user and the generated extended reality environment in a real-time manner. In an example, the electronic device 105 provides a rendering of a scene corresponding to the extended reality environment that can be perceived by the user and interacted with in a real-time manner, such as a host environment for a group session with another user. Additionally, as part of presenting the rendered scene, the electronic device 105 may provide sound, and/or haptic or tactile feedback to the user. The content of a given rendered scene may be dependent on available processing capability, network availability and capacity, available battery power, and current system workload. The electronic device 105 may be, and/or may include all or part of, the electronic system discussed below with respect to FIG. 8.

[0042] The network 106 may communicatively (directly or indirectly) couple, for example, the electronic device 105, the electronic device 110, and/or the electronic device 115 with each other device and/or the server 120. In some examples, the network 106 may be an interconnected network of devices that may include, or may be communicatively coupled to, the Internet.

[0043] The handheld electronic device 104 may be, for example, a smartphone, a portable computing device such as a laptop computer, a companion device (e.g., a digital camera, headphones), a tablet device, a wearable device such as a watch, a band, and the like, or any other appropriate device that includes, for example, one or more speakers, communications

circuitry, processing circuitry, memory, a touchscreen, and/or a touchpad. In some examples, the handheld electronic device 104 may not include a touchscreen but may support touchscreen-like gestures, such as in an extended reality environment. In some examples, the handheld electronic device 104 may include a touchpad. In FIG. 1, by way of example, the handheld electronic device 104 is depicted as a tablet device.

[0044] The electronic device 110 may be, for example, a smartphone, a portable computing device such as a laptop computer, a companion device (e.g., a digital camera, headphones), a tablet device, a wearable device such as a watch, a band, and the like, or any other appropriate device that includes, for example, one or more speakers, communications circuitry, processing circuitry, memory, a touchscreen, and/or a touchpad. In some examples, the electronic device 110 may not include a touchscreen but may support touchscreen-like gestures, such as in an extended reality environment. In some examples, the electronic device 110 may include a touchpad. In FIG. 1, by way of example, the electronic device 110 is depicted as a tablet device. In some examples, the electronic device 110, the handheld electronic device 104, and/or the electronic device 105 may be, and/or may include all or part of, the electronic system discussed below with respect to FIG. 8. In some examples, the electronic device 110 may be another device such as an Internet Protocol (IP) camera, a tablet, or a companion device such as an electronic stylus, etc.

[0045] The electronic device 115 may be, for example, desktop computer, a portable computing device such as a laptop computer, a smartphone, a companion device (e.g., a digital camera, headphones), a tablet device, a wearable device such as a watch, a band, and the like. In FIG. 1, by way of example, the electronic device 115 is depicted as a desktop computer having one or more cameras 150 (e.g., multiple cameras 150). The electronic device 115 may be, and/or may include all or part of, the electronic system discussed below with respect to FIG. 7.

[0046] The server 120 may form all or part of a network of computers or a group of servers 130, such as in a cloud computing or data center implementation. For example, the server 120 stores data and software, and includes specific hardware (e.g., processors, graphics processors and other specialized or custom processors) for rendering and generating content such as graphics, images, video, audio and multi-media files for extended reality environments. In an example, the server 120 may function as a cloud storage server that

stores any of the aforementioned extended reality content generated by the above-discussed devices and/or the server 120.

[0047] FIG. 2 illustrates a block diagram of various components that may be included in electronic device 105, in accordance with aspects of the disclosure. As shown in FIG. 2, electronic device 105 may include one or more cameras such as camera(s) 150 (e.g., multiple cameras 150, each including one or more image sensors 215) that capture images and/or video of the physical environment around the electronic device, one or more sensors 152 that obtain environment information (e.g., depth information) associated with the physical environment around the electronic device 105. Sensors 152 may include depth sensors (e.g., time-of-flight sensors, infrared sensors, radar, sonar, lidar, etc.), one or more microphones, and/or other types of sensors for sensing the physical environment. For example, one or more microphones included in the sensor(s) 152 may be operable to capture audio input from a user of the electronic device 105, such as a voice input corresponding to the user speaking into the microphones. In the example of FIG. 2, electronic device 105 also includes communications circuitry 208 for communication with electronic device 110, electronic device 115, servers 120, and/or other devices and/or systems in some examples. Communications circuitry 208 may include radio frequency (RF) communications circuitry for detecting radio frequency identification (RFID) tags, Bluetooth Low Energy (BLE) communications circuitry, other near-field communications (NFC) circuitry, WiFi communications circuitry, cellular communications circuitry, and/or other wired and/or wireless communications circuitry.

[0048] As shown, electronic device 105 includes processing circuitry 204 (e.g., one or more processors and/or integrated circuits) and memory 206. Memory 206 may store (e.g., temporarily or permanently) content generated by and/or otherwise obtained by electronic device 105. In some operational scenarios, memory 206 may temporarily store images of a physical environment captured by camera(s) 150, depth information corresponding to the images generated, for example, using a depth sensor of sensors 152, meshes and/or textures corresponding to the physical environment, virtual objects such as virtual objects generated by processing circuitry 204 to include virtual content, and/or virtual depth information for the virtual objects. Memory 206 may store (e.g., temporarily or permanently) intermediate images and/or information generated by processing circuitry 204 for combining the image(s) of the physical environment and the virtual objects and/or virtual image(s) to form, e.g.,

composite images for display by display 200, such as by compositing one or more virtual objects onto a pass-through video stream obtained from one or more of the cameras 150.

[0049] As shown, the electronic device 105 may include one or more speakers 211. The speakers may be operable to output audio content, including audio content stored and/or generated at the electronic device 105, and/or audio content received from a remote device or server via the communications circuitry 208.

[0050] Memory 206 may store instructions or code for execution by processing circuitry 204, such as, for example operating system code corresponding to an operating system installed on the electronic device 105, and application code corresponding to one or more applications installed on the electronic device 105. The operating system code and/or the application code, when executed, may correspond to one or more operating system level processes and/or application level processes, such as processes that support capture of images, obtaining and/or processing environmental condition information, and/or determination of inputs to the electronic device 105 and/or outputs (e.g., display content on display 200) from the electronic device 105.

[0051] In some examples, one or more input devices include one or more camera sensors (e.g., one or more optical sensors and/or one or more depth camera sensors such as for tracking a user's gestures (e.g., hand gestures and/or air gestures) as input. In some examples, the one or more input devices are integrated with the computer system. In some examples, the one or more input devices are separate from the computer system. In some examples, an air gesture is a gesture that is detected without the user touching an input element that is part of the device (or independently of an input element that is a part of the device) and is based on detected motion of a portion of the user's body through the air including motion of the user's body relative to an absolute reference (e.g., an angle of the user's arm relative to the ground or a distance of the user's hand relative to the ground), relative to another portion of the user's body (e.g., movement of a hand of the user relative to a shoulder of the user, movement of one hand of the user relative to another hand of the user, and/or movement of a finger of the user relative to another finger or portion of a hand of the user), and/or absolute motion of a portion of the user's body (e.g., a tap gesture that includes movement of a hand in a predetermined pose by a predetermined amount and/or speed, or a shake gesture that includes a predetermined speed or amount of rotation of a portion of the user's body).

[0052] Attention is now directed towards techniques for managing computer-generated environments. Such techniques are described in the context of receiving a request from one or more applications and responding to those requests with information, such as information content the size of a virtual object and/or a pose within an experience of a person. It should be recognized that other configurations can be used with techniques described herein. In addition, techniques optionally complement or replace other techniques for managing computer-generated environments.

[0053] FIG. 3 is a block diagram illustrating a computer system (e.g., computer system 300) in accordance with some examples. Not all of the illustrated components are used in all examples; however, one or more examples can include additional and/or different components than those shown in FIG. 3. In some examples, computer system 300 includes one or more components described above with respect to electronic device 105, handheld electronic device 104, electronic device 110, electronic device 115, and/or server 120 as shown in FIG. 1. Variations in the arrangement and type of the components can be made without departing from the spirit or scope of the claims as set forth herein. Additional components, different components, and/or fewer components can be used as well.

[0054] In some examples, computer system 300 loads, renders, manages, and/or displays computer-generated content in a 3D environment. The 3D environment can be either virtual or physical, with the computer-generated content either completely covering a field of view of a user or supplementing the field of view. For example, computer system 300 can cause a virtual environment to be rendered and displayed to a user such that the user is provided content that is reactive to movements of the user. When the user moves around and performs different gestures, computer system 300 detects and processes the actions to provide tailored information to applications executing on computer system 300.

[0055] As illustrated in FIG. 3, computer system 300 includes 3D environment process 310, 3D framework 320 (e.g., a 3D UI framework and/or other type of 3D framework), 2D framework 330 (e.g., a 2D UI framework and/or other type of 2D framework), display process 340, first user application 350, and second user application 360. While FIG. 3 illustrates that each of these components are on a single computer system, it should be recognized that one or more components can be on another computer system in communication (e.g., wired and/or wireless communication) with computer system 300. In addition, while each component will be discussed separately, in some examples, the

functionality of one or more components are combined together or separated further. In some examples, one or more components of computer system 300 communicate with other components via application programming interfaces (APIs), inter-process communications (IPCs), and/or serial peripheral interfaces (SPIs).

[0056] In some examples, 3D environment process 310 executes as a background process (e.g., a daemon, a service, a system process, an application process, and/or one or more instructions) to manage a 3D environment on behalf of one or more applications (e.g., first user application 350 and/or second user application 360). For example, 3D environment process 310 can create the 3D environment, manage a state of the 3D environment, receive requests from the one or more applications to render content in the 3D environment, communicate with 3D framework 320 and/or 2D framework 330 to service the requests, cause display process 340 to display the 3D environment, and/or detect and process inputs from a number of different sources.

[0057] In some examples, 3D environment process 310 provides one or more APIs to be used by the one or more applications for setting up the 3D environment. In such examples, the APIs can work in a declarative form that allows for developers to create views, animations, and/or other user-interface elements without needing to configure the 3D environment imperatively. In some examples, 3D environment process 310 creates a scene via a scene graph, adds one or more entities to the scene, and/or causes the scene to be rendered.

[0058] In some examples, 3D environment process 310 combines functionality of 3D framework 320 and 2D framework 330 such that user-interface elements and/or functionality provided by 3D framework 320 and/or 2D framework 330 can be used with each other rather than requiring one or the other to be used at a time. For example, 3D environment process 310 acts as a bridge between 3D framework 320 and 2D framework 330, providing each the ability to render objects together in a single scene. In some examples, 3D framework 320 renders 3D objects (e.g., via a first render server) and manages interactions with respect to the 3D objects and/or other objects. Similarly, 2D framework renders 2D objects (e.g., via a second render server different from the first render server) (e.g., and not 3D objects) and manages interactions with respect to the 2D objects and/or other objects. Rather than requiring each framework to work independently, such as providing a separate space for each to own, techniques described herein provide a single space that combines functionality of 3D

framework 320 and 2D framework 330 to create the 3D environment. For example, as further discussed below, 2D environment can render objects to be used by 3D framework 320 when rendering the 3D environment.

[0059] In some examples, to perform such functionality described above, 3D environment process 310 creates a view (e.g., sometimes referred to as a world view) of a 3D environment and adds one or more 3D objects to the view. In such examples, an object of the one or more objects can be hidden, as described further below. In some examples, the object can be used by 3D framework 320 to maintain a place for 2D content from 2D framework 320. In such examples, one technique for implementing such is via a scene graph. The scene graph can include multiple 3D entities that are managed by environment process 310 and/or 3D framework 320. Such 3D entities can include both visible entities and hidden entities. In some examples, a hidden entity (e.g., sometimes referred to as an invisible and/or non-displayed entity) has a size, position, and/or orientation within the 3D environment. Moreover, the hidden entity is connected to a 2D entity such that 3D framework 320 communicates with 2D framework via the hidden entity and/or vice versa.

[0060] FIG. 4 is a block diagram that illustrates an exemplary embodiment for resizing an environment. In some examples, the techniques described below in relation to managing the size of virtual objects in an environment as display of the environment is resized. In some examples, the physical size of a virtual object is maintained, irrespective of whether the angular size of the virtual object is maintained. In some examples, the angular size of the virtual object is maintained, irrespective of the physical size of the virtual object.

[0061] In some examples, the decision to maintain the physical size or angular size of the virtual object is based on one or more characteristics associated with the virtual object. In some examples, if the one or more characteristics (e.g., such as a system flag, designation, type of virtual object, method parameter, etc.) designate for the physical size of the virtual object to be maintained, a computer system would provide information for resizing the virtual object, such that the physical size of the virtual object would be maintained and the angular size of the virtual object would not be maintained. In some examples, if the one or more characteristics designate for the angular size of the virtual object to be maintained, a computer system would provide information (e.g., in response to an API request) for resizing the virtual object, such that the angular size of the virtual object would be maintained and the physical size of the virtual object would not be maintained. In some examples, applications

attempt to maintain the angular size and the physical size of virtual objects while resizing user interfaces and/or environment. However, in some examples, a choice as to whether to prioritize the maintaining of the physical size of the virtual object or the angular size of the virtual object during resizing operations needs to be made. For example, when resizing an object that has text in a virtual environment, a decision can be made to maintain the angular size of the text, irrespective of maintaining the physical size of the text (e.g., and/or the size at which the text was originally displayed and/or would be displayed in the physical environment). In some examples, this decision would allow the text to remain visible to the user as one or more portions in the environment is resized (e.g., get smaller and/or is zoomed out of). For example, when text is included on a billboard and a user's perspective with respect to the billboard becomes further away, the text can maintain the same size as the billboard becomes smaller, allowing the text to maintain its legibility to the user even as the billboard becomes smaller. As another example, when resizing an object that represents a building or structure in the virtual environment, a decision can be made to maintain the physical size of the object in the virtual environment, such that the original size of the virtual object is maintained, irrespective of whether one or more portions are resized (and/or zoomed). For example, the building can include a label that, as the building becomes smaller due to increased distance from a user, the name of the building maintains its angular size with respect to the user and does not become smaller even when the building does not maintain its angular size and becomes smaller.

[0062] FIG. 4 is block diagram 400 that includes object A state block 410 and object B state block 420. Object A state block 410 is a representation of properties that an object (e.g., object A) is tracked (e.g., by a computer system) as having while being presented in an environment. Similarly, object B state block 420 is a representation of properties that a different object is tracked as having while being presented in the environment. In some examples, the objects corresponding to object A state block 410 and object B state block 420 are concurrently displayed and/or presented together in a physical environment. In some examples, the objects (“the respective objects”) corresponding to object A state block 410 and object B state block 420 are hidden in the environment. In some examples, the respective objects are displayed in and/or hidden in content displayed in a mixed- reality and/or virtual- reality environment.

[0063] As represented in object A state block 410 and object B state block 412, object A currently is tracked as having an angular size of one meter and a physical size of two meters while object B is currently tracked as having an angular size of three meters and physical size of one meter. In some examples, object A and/or object B has different sizes than the one represented in object A state block 410 and object B state block 412. In some examples, one or more other properties related to size are tracked for object A and object B. In some examples, one or more properties of other objects are tracked for object A and object B.

[0064] At block 430, a computer system detects that a resize event has occurred. In some examples, the resize event does not include an indication to resize object A or object B. In some examples, it does. In some examples, the resize event is detected by receiving a request from an application. In some examples, the application sends the request via an API call to the computer system to determine the size that object A and/or object B after (or before) the resizing event has occurred.

[0065] Updated object A state block 430 and updated object B state block 450 are provided to show how the resizing event impacted object A and object B differently. Object A state block 430 indicates that object A is currently being tracked to have an angular size of one meter and a physical size of five meters. Thus, at block 430, object A's angular size has been maintained before and after the occurrence of the resizing event. However, object B's angular size has not been maintained. In fact, object B state block 412 object B's angular size has increased from three meters to five meters while the physical size of object B has remained the same at one meter (e.g., when comparing object B state block 420 to updated object B state block 450). In some examples, a computer system sends a response to an API call that indicates how object A and object B should be resized after the occurrence of the resizing event.

[0066] It should be understood that the scenario described above in relation to resizing objects in different manners and/or tracking can occur in many different user interfaces and/or environments. In some examples, the size of object A and object B are tracked while each of the objects are displayed in an application window, such as a two-dimensional window that is displayed on a surface or a three-dimensional window that is bounded. In some examples, whether the angular size or physical size is maintained is determined based on a particular mode in which a computer system is operating. For example, the physical size of a virtual object can be maintained while the computer system is operating in a two-

dimensional and/or a window mode, where this paradigm can be switched (e.g., angular size maintained) after the computer system shift to operating in a three-dimensional and/or an immersive mode.

[0067] FIG. 5 is a flow diagram illustrating a method (e.g., method 500) for managing the size of virtual objects in accordance with some examples. Some operations in method 500 are, optionally, combined, the orders of some operations are, optionally, changed, and some operations are, optionally, omitted.

[0068] As described below, method 500 provides an intuitive way for managing the size of virtual objects. Method 500 reduces the cognitive burden on a user for managing an experience, thereby creating a more efficient human-machine interface. For battery-operated computing devices, enabling a user to manage the size of virtual objects faster and more efficiently conserves power and increases the time between battery charges.

[0069] In some examples, method 500 is performed by and/or at a system process (e.g., a daemon, a service, and/or other type of system process) (e.g., a process of an operating system) of a computer system (e.g., a device, a personal device, a user device, and/or a head-mounted display (HMD)). In some examples, the computer system is in communication with input/output devices, such as one or more cameras (e.g., a telephoto camera, a wide-angle camera, and/or an ultra-wide-angle camera), speakers, microphones, sensors (e.g., heart rate sensor, monitors, antennas (e.g., using Bluetooth and/or Wi-Fi), fitness tracking devices (e.g., a smart watch and/or a smart ring), and/or near-field communication sensors). In some examples, the computer system is in communication with a display generation component (e.g., a projector, a display, a display screen, a touch-sensitive display, and/or a transparent display).

[0070] At 502, the computer system receives, from a first application (e.g., a user application and/or an application installed on the computer system (e.g., by a user and/or another computer system)) (e.g., of the computer system), a request to render a first object (e.g., corresponding to the application) in an environment (e.g., multi-dimensional environment, such as a two-dimensional or a three-dimensional environment) (e.g., a virtual or a physical environment). In some examples, the computer system is a phone, a watch, a tablet, a fitness tracking device, a wearable device, a television, a multi-media device, an accessory, a speaker, and/or a personal computing device.

[0071] At 504, in response to receiving the request to render the first object in the environment, the computer system renders the first object.

[0072] At 506, after rendering the first object and while the first object is being displayed in the environment, the computer system detects an indication of (e.g., via one or more input devices (e.g., as described above)) (e.g., an input (e.g., a tap input and/or a non-tap input, such as an air input (e.g., a pointing air gesture, a tapping air gesture, a swiping air gesture, and/or a moving air gesture), a gaze input, a gaze-and-hold input, a mouse click, a mouse click-and-drag, a key input of a keyboard, a voice command, a selection input, and/or an input that moves the computer system in a particular direction and/or to a particular location) corresponding to, data corresponding to, and/or a call via one or more APIs) a request to change a size of content provided by the first application (e.g., a request to change a window size and/or a zoom size and/or a request to change a position of a view of the environment from a first position (e.g., a first location within the environment) to a second position (e.g., a second location within the environment) different from the first position).

[0073] In response to detecting the indication of the request to change the size of content provided by the first application, rendering a representation of the environment by: in accordance with (at 508) a determination that the environment is rendered according to (e.g., instructed to be displayed as being in and/or operating in, displayed in, and/or is in) a first mode (e.g., an application mode, such as a window mode and/or a mode that has greater restraints on size of objects than the second mode described below) and that the first object is a first type of object (e.g., a true scale object), wherein the environment corresponds to the first application and a second application different from the first application (and/or one or more other applications) while the environment is rendered according to the first mode, maintaining a physical size (e.g., in real world units) (e.g., within a predefined margin) of the first object without maintaining an angular size (e.g., relative to user's viewpoint) of the first object; in accordance with (at 510) a determination that the environment is rendered according to the first mode and that the first object is a second type of object (e.g., dynamically scale object) different from the first type of object, maintaining the angular size of the first object without maintaining the physical size of the first object; and in accordance with (at 512) a determination that the environment is rendered according to a second mode (e.g., an application mode, such as an immersive mode) (e.g., without respect to whether the first object is the first type of object or the second type of object), wherein the environment

corresponds to the first application without corresponding to another application (e.g., another user application and the second application) (e.g., not including a system process and/or application) different from the first application (e.g., the environment includes one or more entities corresponding to the first application without including an entity corresponding to the other application) while the environment is rendered according to the second mode, maintaining the physical size of the first object without maintaining the angular size of the first object. In some examples, maintaining the physical size includes providing the physical size. In some examples, maintaining the angular size includes providing the angular size.

[0074] In some examples, the second mode is a window mode. In some examples, the first object is within a window (e.g., a two-dimensional or three-dimensional window) (e.g., a graphical user interface, an application window, a window suspended in space, a window that is surrounded by a three-dimensional environment, and/or a window that is displayed concurrently with one or more other windows) (e.g., a separate viewing area) while the environment is rendered according to the second mode.

[0075] In some examples, while a second object is being displayed in the environment, the computer system detects a second indication of (e.g., via one or more input devices (e.g., as described above)) (e.g., an input (e.g., a tap input and/or a non-tap input, such as an air input (e.g., a pointing air gesture, a tapping air gesture, a swiping air gesture, and/or a moving air gesture), a gaze input, a gaze-and-hold input, a mouse click, a mouse click-and-drag, a key input of a keyboard, a voice command, a selection input, and/or an input that moves the computer system in a particular direction and/or to a particular location) corresponding to, data corresponding to, and/or a call via one or more APIs) a request to change a size of content (e.g., provided by the first application (e.g., a request to change a window size and/or a zoom size and/or a request to change a position of a view of the environment from a first position (e.g., a first location within the environment) to a second position (e.g., a second location within the environment) different from the first position)), wherein the second indication is different from the indication. In some examples, the second object is provided by the first application. In some examples, the second indication is separate from the first indication. In some examples, in response to detecting the second indication of the request to change the size of content, the computer system renders a second representation (e.g., the representation and/or another representation) of the environment by: in accordance with a determination that the environment is rendered according to the second mode and that the

second object is a third type of object (e.g., the first type of object or another type of object different from the first type of object), the computer system maintains a physical size (e.g., in real world units) (e.g., within a predefined margin) of the second object without maintaining an angular size (e.g., relative to user's viewpoint) of the second object. In some examples, maintaining the physical size of the second object includes sending (e.g., to the first application) the physical size of the second object, where the physical size of the second object does not change as the size of content is increased and/or decreased. In some examples, not maintaining the angular size of the second object includes sending (e.g., to the first application) the angular size of the second object, where the angular size of the second object changes as the size of content is increased and/or decreased; and in accordance with a determination that the environment is rendered according to the second mode and that the second object is a fourth type of object (e.g., the second type of object or another type of object different from the second type of object) different from the first type of object, the computer system maintains the angular size of the first object without maintaining the physical size of the first object, wherein the fourth type of object is different from the third type of object. In some examples, not maintaining the physical size of the second object includes sending the physical size of the second object, where the physical size of the second object does not change as the size of content is increased and/or decreased. In some examples, maintaining the angular size of the second object includes sending the angular size of the second object, where the angular size of the second object changes as the size of content is increased and/or decreased.

[0076] In some examples, while a third object is being displayed concurrently with the first object in the environment and in response to detecting the indication of the request to change the size of content provided by the first application, rendering the representation of the environment by: while maintaining the angular size of the first object (e.g., and in accordance with a determination that the environment is rendered according to the first mode and that the third object is the first type of object), maintaining a physical size (e.g., in real world units) (e.g., within a predefined margin) of the third object without maintaining an angular size (e.g., relative to user's viewpoint) of the third object; and while maintaining the physical size of the first object (e.g., and in accordance with a determination that the environment is rendered according to the first mode and that the second object is the second type of object), maintaining the angular size of the third object without maintaining the physical size of the third object. In some examples, while the third object is being displayed

concurrently with the first object in the environment, in response to detecting the indication of the request to change the size of content provided by the first application, and in accordance with a determination that the environment is rendered according to the second mode, the computer system renders the representation of the environment by maintaining the angular size of the third object without maintaining the physical size of the third object.

[0077] In some examples, while a fourth object of a second application is being displayed concurrently with the first object in the environment and in response to detecting an indication of a request to change a size of content (e.g., provided by the first application and/or the second application), wherein the second application is different from the first application, the computer system renders the representation of the environment by: while maintaining the angular size of the first object (e.g., and in accordance with a determination that the environment is rendered according to the first mode and that the fourth object is the first type of object), the computer system maintains a physical size (e.g., in real world units) (e.g., within a predefined margin) of the fourth object without maintaining an angular size (e.g., relative to user's viewpoint) of the fourth object; and while maintaining the physical size of the first object (e.g., and in accordance with a determination that the environment is rendered according to the first mode and that the fourth object is the second type of object), the computer system maintains the angular size of the fourth object without maintaining the physical size of the fourth object.

[0078] In some examples, the content provided by the first application is concurrently displayed in the environment with content provided by the second application (e.g., while the environment is rendered according to the first mode and, in some examples, while the environment is not operating in the second mode).

[0079] In some examples, while the environment is rendered according to the first mode, the computer system detects a request (e.g., via one or more input devices (e.g., as described above)) (e.g., an input (e.g., a tap input and/or a non-tap input, such as an air input (e.g., a pointing air gesture, a tapping air gesture, a swiping air gesture, and/or a moving air gesture), a gaze input, a gaze-and-hold input, a mouse click, a mouse click-and-drag, a key input of a keyboard, a voice command, a selection input, and/or an input that moves the computer system in a particular direction and/or to a particular location) corresponding to, data corresponding to, and/or a call via one or more APIs) to transition to the second mode. In some examples, in response to detecting the request to transition to the second mode, the

computer system causes the environment to operate in the second mode, wherein causing the environment to operate in the second mode includes removing, from the environment, content provided by another application different from the first application. In some examples, in response to detecting the request to transition to the second mode while and/or after maintaining the physical size of the first object without maintaining the angular size of the first object, the computer system continues to maintain the physical size of the first object without maintaining the angular size of the first object while the environment is rendered according to the second mode. In some examples, in response to detecting the request to transition to the second mode while and/or after maintaining the angular size of the first object without maintaining the physical size of the first object, the computer system continues to maintain the angular size of the first object without maintaining the physical size of the first object while the environment is rendered according to the second mode.

[0080] In some examples, while the environment is rendered according to the first mode: content corresponding to the first application is in a first area of the environment and not in a second area of the environment and content corresponding to a second application is in the first area and not in the second area, wherein the second area is different from the first area, and wherein the second application is different from the first application. In some examples, the second area is separate from and/or not surrounded by the first area, and the first area is separated from and/or not surrounded by the second area.

[0081] In some examples, while the environment is rendered according to the first mode and in accordance with a determination that a first set of one or more criteria is satisfied (e.g., that a current window corresponding to the first application is a three-dimensional bounded window), content provided by the first application is within a three-dimensional bounded window (e.g., a graphical user interface) (e.g., a separate viewing area). In some examples, content within a three-dimensional bounded window can move a distance in the x, y, and/or z directions, content within a three-dimensional bounded window can move a distance in each of the x, y, and z directions but is limited to moving a certain amount (e.g., non-zero) in each of the x, y, and/or z directions and/or within a three-dimensional bounded window is displayed with depth.

[0082] In some examples, while in the first mode and in accordance with a determination that a second set of one or more criteria is satisfied (e.g., that a current window corresponding to the first application is a two-dimensional bounded window), content corresponding to the

first application is within a two-dimensional bounded window (e.g., a graphical user interface) (e.g., a separate viewing area). In some examples, the second set of one or more criteria is different from the first set of one or more criteria. In some examples, content within a two-dimensional bounded window can move a distance in the x and y directed but cannot move in the z direction, content within a three-dimensional bounded window can move a distance in each of the x and y directions but is limited to moving a certain amount (e.g., non-zero) in each of the x, y, and z directions (e.g., where content cannot move in the z direction), and/or within a three-dimensional bounded window is not displayed with depth.

[0083] In some examples, the first set of one or more criteria includes a criterion that is satisfied when a first window (e.g., the three-dimensional bounded window) corresponding to the first application includes a style (e.g., a setting, an option, a property, and/or data included in data of the first window) defining that the window is a three-dimensional bounded window. In some examples, the second set of one or more criteria includes a criterion that is satisfied when a second window (e.g., the two-dimensional bounded window) corresponding to the first application includes a style defining that the second window is a two-dimensional bounded window.

[0084] In some examples, the second area surrounds the first area. In some examples, the first area surrounds the second area.

[0085] In some examples, maintaining the physical size of the first object without maintaining the angular size of the first object includes, at a first time while the size of content provided by the first object is changed, maintaining the physical size of the first object and maintaining the angular size of the first object.

[0086] In some examples, maintaining the physical size of the first object without maintaining the angular size of the first object includes: at a second time, different from the first time, while the size of content provided by the first object is changed, changing the angular size of the first object while maintaining the physical size of the first object. In some examples, maintaining the angular size of the first object without maintaining the physical size of the first object includes: at a third time while the size of content provided by the first object is changed maintaining the physical size of the first object and maintaining the angular size of the first object; and at a fourth time, different from the third time, while the size of

content provided by the first object is changed, changing the physical size of the first object while maintaining the angular size of the first object.

[0087] Note that details of the processes described above with respect to method 500 (e.g., FIG. 5) are also applicable in an analogous manner to other methods described herein. For example, method 700 optionally includes one or more of the characteristics of the various methods described above with reference to method 500. For example, the size of virtual objects can be managed using one or more steps of method 500 while a location of a user can be determined using one or more steps of method 700 for the same experience and/or at the same time. For brevity, these details are not repeated below.

[0088] FIGS. 6A-6B are exemplary user interfaces for managing an experience that is attended by one or more users. The user interfaces of FIGS. 6A-6B are used to illustrate the processes described below, including the processes of FIG. 7.

[0089] FIG. 6A illustrates display area 600 including origin 610 (e.g., the “X”) and person 620. Origin 610 is the origin (e.g., the center, the meeting point, the target area, and/or the performing place) for experience 640. In some examples, an experience includes a meeting, game, an event, and/or a conference that is executing in a virtual environment. In some examples, the origin for an experience is the focal point of the experience, such as bowling position at a bowling alley, a stage at a concert, the podium at a conference, and/or the batter’s box at a baseball game. It should be understood that display area 600 is a display area a computer system, such as an HMD and/or smart glasses. In some examples, the display area is a display generation component, such as a touch sensitive display, a translucent display, and/or a display screen. In some examples, the computer system includes one or more components as those discussed above in relation to devices 105, 110, and/or 300.

[0090] At FIG. 6A, the display area includes an indication that person 620 is located at the origin. Before FIG. 6A, a determination was made that only person 620 (or less than a number of people) would be attending experience 640 at a particular point in time. Thus, because this determination was made, person 620 is located at the origin in FIG. 6A. In some examples, the determination was made at a time that experience 640 was launched and/or one or a user interface corresponding experience 640 was initially displayed in the virtual environment. In some examples, the experience represented in FIGS. 6A-6B was launched

from another experience in the virtual environment that included person 620 or person 630 of FIG. 6B.

[0091] FIG. 6B illustrates display area 600 including person 620 and person 630 placed at location outside of origin 610. Here, person 620 and person 630 are placed outside of the origin based on a determination that more than a threshold number of people will be attending the experience. In some examples, the determination was made at a time that experience 640 was launch and/or a user interface corresponding experience 640 was displayed in the virtual environment.

[0092] FIG. 7 is a flow diagram illustrating a method (e.g., method 700) for managing an experience in accordance with some examples. Some operations in method 700 are, optionally, combined, the orders of some operations are, optionally, changed, and some operations are, optionally, omitted.

[0093] As described below, method 700 provides an intuitive way for managing an experience. Method 700 reduces the cognitive burden on a user for managing an experience, thereby creating a more efficient human-machine interface. For battery-operated computing devices, enabling a user to manage an experience faster and more efficiently conserves power and increases the time between battery charges.

[0094] At 702, the computer system sends, from an application (e.g., the application described above with respect to method 700) (and, in some examples, via an application programming interface (API)), a request for a pose (e.g., from the 610 relative to 620 and 630) (e.g., a location and/or an orientation) associated with (and/or corresponding to) a user (e.g., a person being tracked by the application and/or the computer system) with respect to an origin (e.g., 610) (e.g., a coordinate location, a middle, a center, and/or a designed origin) of an experience (e.g., 640) (e.g., a game, a central experience, and/or meeting point (e.g., socializing point and/or origin), and/or a currently viewable portion) for the application.

[0095] At 704, after (e.g., as a response to) sending the request for the pose associated with the user with respect to the origin of the experience for the application and in accordance with a determination that a first number of users (e.g., one or, in some examples, more than one) are in the experience (e.g., a non-shared experience and/or an individual experience) for

the application, the computer system receives (e.g., via the API) a first pose (e.g., as described above in relation to FIG. 6A).

[0096] At 706, after sending the request for the pose associated with the user with respect to the origin of the experience for the application and in accordance with the determination that the first number of users are in the experience for the application, the computer system renders the experience for the application based on (e.g., from, in a pose that is the same as, opposite of, and/or adjusted using and/or from) the first pose (e.g., as described above in relation to FIG. 6A).

[0097] At 708, after sending the request for the pose associated with the user with respect to the origin of the experience for the application and in accordance with a determination that a second number of users (e.g., more than one or, in some examples, one) are in the experience (e.g., a shared experience and/or an experience involving a plurality of users) for the application, the computer system receives (e.g., via the API) a second pose different from the first pose (e.g., as described above in relation to FIG. 6B).

[0098] At 710, after sending the request for the pose associated with the user with respect to the origin of the experience for the application and in accordance with a determination that a second number of users (e.g., more than one or, in some examples, one) are in the experience (e.g., a shared experience and/or an experience involving a plurality of users) for the application, the computer system renders the experience for the application based on (e.g., from, in a pose that is the same as, opposite of, and/or adjusted using and/or from) the second pose e.g., as described above in relation to FIG. 6B). In some examples, after sending the request for the pose associated with the user with respect to the origin of the experience for the application and in accordance with a determination that the second number of users are in the experience for the application, the computer system does not receive the first pose. In some examples, after sending the request for the pose associated with the user with respect to the origin of the experience for the application and in accordance with a determination that the second number of users are in the experience for the application, the computer system does not render the experience for the application based on the first pose. In some examples, after sending the request for the pose associated with the user with respect to the origin of the experience for the application and in accordance with a determination that the first number of users are in the experience for the application, the computer system does not receive the second pose. In some examples, after sending the request for the pose associated with the

user with respect to the origin of the experience for the application and in accordance with a determination that the first number of users are in the experience for the application, the computer system does not render the experience of the application based on the second pose.

[0099] In some examples, the first number of users is one. In some examples, the first pose corresponds to a center (e.g., a center of origin) of the experience for the application. In some examples, the second number of users is two or more. In some examples, the second pose corresponds to an offset greater than zero from the center of the experience for the application. In some examples, rendering the experience for the application based on the first pose includes rendering the experience from a perspective located at the center of the experience for the application. In some examples, rendering the experience for the application based on the second pose includes rendering the experience from a perspective located at the offset from the center of the experience for the application.

[0100] In some examples, after (e.g., as a response to and/or in response to) sending the request for the pose associated with the user with respect to the origin of the experience for the application and in accordance with a determination that a third number of users (e.g., more than two) are in the experience for the application, wherein the third number of users is different from the first number of users and the second number of users, the computer system receives (e.g., via the API) a third pose different from the first pose and the second pose. In some examples, after (e.g., as a response to and/or in response to) sending the request for the pose associated with the user with respect to the origin of the experience for the application and in accordance with a determination that a third number of users (e.g., more than two) are in the experience for the application, wherein the third number of users is different from the first number of users and the second number of users, the computer system renders the experience for the application based on the third pose.

[0101] In some examples, the second pose is generated based on a plurality of users being placed around an origin (e.g., a center and/or center of origin) of the experience for the application. In some examples, placing the plurality of users around the origin includes not placing a user at the origin.

[0102] In some examples, the second pose is generated based on multiple users being placed in a sequence (e.g., each user next to and/or adjacent to another user and/or each user in a line, circle, and/or other geometric pattern) near (e.g., around, a non-zero distance away

from, and/or within a predefined distance from) a second origin (e.g., a center) of the experience for the application. In some examples, placing the plurality of users in the sequence near the origin includes not placing a user at the origin.

[0103] In some examples, the request for the pose is sent in response to detecting intent of a user to launch (e.g., initiate and/or start) the application (and/or detecting launch of the application and/or detecting a request to launch the application and/or before detecting the request to launch the application).

[0104] In some examples, after (and/or while) detecting intent of the user to launch the application in a respective mode (e.g., the second mode as described above with respect to method 500), a process of a computer system executing the application performs one or more pre-launch tasks (e.g., causing one or more processes (e.g., non-essential processes) to shut down and/or cease executing) (e.g., configuring one or more states and/or modes of the computer system (e.g., an immersion state and/or the second mode) and/or transitioning the computer system from operating in the first mode, as described above with respect to method 500, to operating in the second mode, as described above with respect to method 500) on behalf of the application. In some examples, the one or more pre-launch tasks correspond to the application launching in the respective mode. In some examples, the one or more pre-launch tasks occur before an input is detected to launch the application in the respective mode. In some examples, detecting the intent of the user includes detecting that a user is going to and/or is likely going to send a request (e.g., moving gaze towards a button, starting to issue a voice command, starting to perform one or more air gestures, and/or moving towards a virtual object corresponding to the application) to launch the application (e.g., within a predetermined period of time (e.g., 1-60 seconds)).

[0105] In some examples, the request for the pose is sent in conjunction with a request to change from a second experience to the experience (e.g., while participating in the second experience) (e.g., going from the second experience to the experience) (e.g., going from virtual a conference room to a virtual bowling game). In some examples, the second experience is for the application. In some examples, the second experience is different from the experience.

[0106] In some examples, the first pose includes a first displacement (e.g., zero or more units of distance) from a third origin (e.g., a center and/or center of origin) of the experience.

In some examples, the second pose includes a second displacement (e.g., zero or more units of distance) from the third origin of the experience. In some examples, the second displacement is different from (e.g., more than and/or less than) the first displacement.

[0107] In some examples, the request for the pose includes an indication of a coordinate space (e.g., a type of coordinate space, such as relative to a viewpoint, a location in an environment, and/or a fourth origin of the experience). In some examples, the request for the pose includes a request to identify a coordinate space.

[0108] In some examples, in accordance with a determination that the request for the pose includes an indication of a first coordinate space (e.g., a type of coordinate space, such as relative to a viewpoint, a location in an environment, and/or a fourth origin of the experience), a respective pose (e.g., the first pose and/or the second pose) that is received corresponds to the first coordinate space (e.g., a cartesian and/or a Polar coordinate system). In some examples, in accordance with a determination that the request for the pose includes an indication of a second coordinate space (e.g., a cartesian and/or a Polar coordinate system) different from the first coordinate space, the respective pose that is received corresponds to the second coordinate space.

[0109] In some examples, after receiving the first pose, the computer system sends a second request for the pose associated with the user (e.g., with respect to the origin of the experience for the application). In some examples, after sending the second request for the pose associated with the user, the computer system receives one or more affine transformations instead of a pose, wherein the one or more affine transformations correspond to a previous pose (e.g., the first pose, the second pose, or another pose different from the first pose and the second pose) received by the application. In some examples, the one or more affine transformations are received instead and/or in lieu of the first pose.

[0110] Note that details of the processes described above with respect to method 700 (e.g., FIG. 7) are also applicable in an analogous manner to the methods described herein. For example, method 500 optionally includes one or more of the characteristics of the various methods described above with reference to method 700. For brevity, these details are not repeated below.

[0111] Applications in mixed reality operating systems have scenes that have an extent in the physical space around the user. In some examples, two-dimensional framework measures sizes in points. In some examples, in mixed reality Operating System 1.0, certain kinds of scenes are scaled dynamically depending on where the user places them in the environment, causing points' physical length to increase or decrease accordingly. There are many use cases for letting an application express certain lengths in real-world units: to display UI that is size-invariant, to initially size volumetric windows, to ensure certain UI elements can be replaced if they would be too small to see, to overlay 3D UI on the outside world, and so on. We propose API for these use cases.

[0112] The new API entry points provided cover the following: Scene sizing: Scenes whose content should match real-world measurements can specify their default size in those measurements, rather than points. This provides guidance to the system as to how to scale and size them on first appearance. (This is the only entry point that takes real-world measurements instead of points.) Points-length conversions for arbitrary values: Since the points-per-meter value depends on the particular scene (both its type and in some cases its positioning), the conversion machinery ties into the two-dimensional framework view update system and is only available from inside a View's implementation. These API introduce the assumption that points and physical sizes can relate in well-defined, device-detectable ways. The variety of output mechanisms of devices running other operating systems makes it fraught to make the same guarantee outside of mixed reality operating system. In some examples, the API is available on mixed reality operating system.

[0113] The following definitions make use of the UnitLength type vended by the Foundation framework. The following modifiers are added as an extension to Scene. The documentation comment is written once, but in the final result it will be repeated, appropriately slightly edited, for each new variant.

[0114] ```swift`

[0115] `@available(phone operating systems, unavailable)`

[0116] `@available(computer operating system, unavailable)`

[0117] `@available(media operating system, unavailable)`

[0118] @available(wearable operating system , unavailable)

[0119] @available(mixed reality operating system 1.0, *)

[0120] extension Scene {

[0121]

[0122] /// Sets a default width and height for a window with an extent in space.

[0123] ///

[0124] /// Use this scene modifier to indicate a default initial size for a new

[0125] /// window that the system creates from a ``Scene`` declaration. For

[0126] /// example, a user can request that new windows that a ``WindowGroup``

[0127] /// generates occupy 30 cm in the x-dimension and 60 cm in

[0128] /// the y-dimension:

[0129] ///

[0130] /// @main

[0131] /// struct MyApp: App {

[0132] /// var body: some Scene {

[0133] /// WindowGroup {

[0134] /// ContentView()

[0135] /// }

[0136] /// .windowStyle(.volumetric)

[0137] /// .defaultSize(width: 0.3, height: 0.6, in: .meters)

[0138] /// }

[0139] /// }

[0140] ///

[0141] /// The size that a user specifies acts only as a default for when the window

[0142] /// first appears. If appropriate, users can later resize the window using

[0143] /// interface controls that the system provides. Also, during state restoration,

[0144] /// the system restores windows to their most recent size rather than

[0145] /// the default size.

[0146] ///

[0147] /// In some examples, if a user specifies a default size that's outside the range of
the window's

[0148] /// inherent resizability in one or both dimensions, the system clamps the

[0149] /// affected dimension to keep it in range. A user can configure the

[0150] /// resizability of a scene using the ``Scene/windowResizability(\_):``

[0151] /// modifier.

[0152] ///

[0153] /// In some examples, this default size modifier affects any scene type whose
window style is

[0154] /// ``WindowStyle/volumetric``. It has no effect on non-volumetric scenes.

[0155] ///

[0156] /// If a user wants to specify the size input in terms of a size instance,

[0157] /// use ``Scene/defaultSize(\_):`` instead.

[0158] ///

[0159] /// - Parameter width: The default width for windows created from a scene.

[0160] /// - Parameter height: The default height for windows created from a scene.

[0161] /// - Parameter unit: The measurement unit of physical extent that `width` and

[0162] /// `height` are specified in.

[0163] ///

[0164] /// - Returns: A scene that uses a default size for new windows.

[0165] public func defaultSize(
 [0166] width: CGFloat,
 [0167] height: CGFloat,
 [0168] depth: CGFloat,
 [0169] in unit: UnitLength) -> some Scene

[0170]

[0171] /// Same doc comment as above, with minor appropriate edits.

[0172] public func defaultSize(
 [0173] _ size: Size3D,
 [0174] in unit: UnitLength) -> some Scene

[0175] }

[0176] ``

[0177] In some examples, a modifier request that a window start with the specified width, height and depth. In some examples, in mixed reality operating system 1.0, this modifier is respected for scenes that respect the .volumetric window style. For example:

[0178] ```swift

```

[0179] struct MyAquariumApp: App {
[0180]     var body: some Scene {
[0181]         WindowGroup(id: "aquarium") {
[0182]             FishTankView()
[0183]         }
[0184]         .windowStyle(.volumetric)
[0185]         .defaultSize(width: 0.9, height: 0.45, depth: 0.4, in: .meters)
[0186]     }
[0187] }
[0188] ```

```

[0189] This requests that the volume be of 90 cm by 45 cm by 40 cm. In some examples, since volumetric scenes are statically sized in mixed reality operating system 1.0, this also sets the scene's size for the whole of the application lifetime.

[0190] Two APIs are proposed: a conversion type that can be accessed in the environment, and a convenience type that simplifies use in the case of constant real-world length values.

```

[0191] ```swift
[0192] extension EnvironmentValues {
[0193]     /// The physical metrics associated with a scene.
[0194]     ///
[0195]     /// Reading this value returns a `PhysicalMetricsConverter` corresponding
[0196]     /// to the window scene associated with the environment's reader. In some
    examples, the

```

[0197] /// converter can convert point sizes into physical measurements of length,

[0198] /// and vice versa.

[0199] ///

[0200] /// In some examples, reading this value is supported in the body of a ``View``
or of

[0201] /// a type that inherits a ``View``'s environment.

[0202] @available(phone operating system, unavailable)

[0203] @available(computer operating system, unavailable)

[0204] @available(media operating system, unavailable)

[0205] @available(wearable operating system , unavailable)

[0206] @available(mixed reality operating system 1.0, *)

[0207] public var physicalMetrics: PhysicalMetricsConverter { get }

[0208] }

[0209]

[0210] /// In some examples, a physical metrics converter provides conversion between
point values and

[0211] /// their extent in 3D space, in the form of physical length measurements.

[0212] ///

[0213] /// In some examples, converters are available from the environment of a `View`
or other type that

[0214] /// inherits a `View`'s environments, and are associated with the scene that

[0215] /// contains that environment. The conversions expect (or emit) values in points

[0216] /// in that scene's coordinate system, and convert these to (or from) measurements

[0217] /// of length in the user's reference frame. For example, if the scene is

[0218] /// scaled, that scale will be taken into account.

[0219] ///

[0220] /// To obtain a converter, use the ``EnvironmentValues/physicalMetrics`` key:

[0221] ///

[0222] /// ```swift

[0223] /// struct MyView: View {

[0224] /// @Environment(\.physicalMetrics) var physicalMetrics

[0225] /// ...

[0226] /// }

[0227] /// ```

[0228] ///

[0229] /// In some examples, when user action modifies a scene so that measurements have changed (e.g., by

[0230] /// changing its scale), the view that accessed that environment key and its hierarchy

[0231] /// will be reevaluated.

[0232] ///

[0233] /// Attempting to obtain a converter inside a type not associated with a scene's

[0234] /// contents (for example, from an ``App`` or ``Scene``'s environment) is not

[0235] /// supported.

[0236] @available(phone operating system, unavailable)

[0237] @available(computer operating system, unavailable)

[0238] @available(media operating system, unavailable)

[0239] @available(wearable operating system, unavailable)

[0240] @available(mixed reality operating system 1.0, *)

[0241] public struct PhysicalMetricsConverter {

[0242]

[0243] /// Converts a vector of physical length measurements, in the specified unit,

[0244] /// to a vector of values in points suitable for use in the environment this

[0245] /// converter is associated with.

[0246] ///

[0247] /// - Parameters:

[0248] /// - `lengths`: A vector of physical measurements of length

[0249] /// - `unit`: The unit of measure for the lengths

[0250] /// - Returns: A value in points. Use this value only in the scene this

[0251] /// converter was associated with.

[0252] public func convert<V>(_ lengthValues: V, from unit: UnitLength) -> V where
V : VectorArithmetic

[0253]

[0254] /// Converts a vector of values in points to corresponding physical length

[0255] /// measurements in the specified unit,

[0256] ///

[0257] /// In some examples, the point values are in the coordinate system of the scene

[0258] /// that this converter is associated with. In some examples, if the scene is scaled, the

[0259] /// physical measurement will take this scale into account.

[0260] ///

[0261] /// - Parameters:

[0262] /// - `points`: A vector of values in points in the associated scene's

[0263] /// coordinate system

[0264] /// - `unit`: The unit of measure for the returned lengths

[0265] /// - Returns: A vector of physical length measurements, each converted from the

[0266] /// points value in the input vector at the same position.

[0267] /// converter was associated with.

[0268] public func convert<V>(_ pointValues: V, to unit: UnitLength) -> V where V : VectorArithmetic

[0269]

[0270] /// Same doc comment, with appropriate minor edits, for all the following:

[0271] public func convert(_ point: Point3D, from unit: UnitLength) -> Point3D

[0272] public func convert(_ point: Point3D, to unit: UnitLength) -> Point3D

[0273] public func convert(_ size: Size3D, from unit: UnitLength) -> Size3D

[0274] public func convert(_ size: Size3D, to unit: UnitLength) -> Size3D

[0275] public func convert(_ rect: Rect3D, from unit: UnitLength) -> Rect3D

```

[0276]     public func convert(_ rect: Rect3D, to unit: UnitLength) -> Rect3D

[0277]     public func convert(_ point: CGPoint, from unit: UnitLength) -> CGPoint

[0278]     public func convert(_ point: CGPoint, to unit: UnitLength) -> CGPoint

[0279]     public func convert(_ size: CGSize, from unit: UnitLength) -> CGSize

[0280]     public func convert(_ size: CGSize, to unit: UnitLength) -> CGSize

[0281]     public func convert(_ rect: CGRect, from unit: UnitLength) -> CGRect

[0282]     public func convert(_ rect: CGRect, to unit: UnitLength) -> CGRect

[0283]     public func convert(_ lengthValue: CGFloat, from unit: UnitLength) ->
CGFloat

[0284]     public func convert(_ pointsValue: CGFloat, to unit: UnitLength) -> CGFloat

[0285]

[0286]     }

[0287]     """The PhysicalMetricsConverter type provides points-to-lengths conversions, and
vice versa, for several types. There are no public constructors for this type. It is accessed
from the environment using the \.physicalMetrics key:

[0288]     """

[0289]     struct RulerView: View {

[0290]         @Environment(\.physicalMetrics) var metrics: PhysicalMetricsConverter

[0291]

[0292]         let size: CGFloat

[0293]

[0294]         var lengthInPoints: CGFloat {

```

[0295] metrics.convert(size, from: .meters)

[0296] }

[0297] }

[0298] ``

[0299] Using this environment key outside of a View that is in a view hierarchy is not supported; see below. In some examples, implementations that perform a conversion will form a dependency on the scene's points-per-meter value, and will be reevaluated as the points-per-meter for a scene changes (for example, a 2.5D window being repositioned by the user). Using this environment key outside of a View that is in a view hierarchy is not supported; see below. In some examples, implementations that perform a conversion will form a dependency on the scene's points-per-meter value, and will be reevaluated as the points-per-meter for a scene changes (for example, a 2.5D window being repositioned by the user).

[0300] In some examples, if a view makes use of fixed real-world length measurements, it can define a name for them using the PhysicalMetric type. This is similar to the existing ScaledMetric type already in two-dimensional framework; unlike that type. In some examples, the metric isn't scaled from points to points, but from a real-world measurement to points.

[0301] ``swift

[0302] /// Provides access to a value in points that corresponds to the specified

[0303] /// physical measurement.

[0304] ///

[0305] /// Use this property wrapper inside a ``View`` or a type that inherits a ``View``'s

[0306] /// environment, like a ``ViewModifier``. Its value will be the equivalent in points

[0307] /// of the physical measurement of length you specify.

[0308] ///

[0309] /// For example, to have a variable that contains the amount of points
corresponding

[0310] /// to one meter, you can do the following:

[0311] ///

[0312] /// ```swift

[0313] /// struct MyView: View {

[0314] /// @PhysicalMetric(1.0, in: .meters) var oneMeter

[0315] /// ...

[0316] /// }

[0317] /// ```

[0318] ///

[0319] /// Using this wrapper for a property of a type not associated with a scene's view

[0320] /// contents, like an ``App`` or a ``Scene``, is unsupported.

[0321] @available(phone operating system, unavailable)

[0322] @available(computer operating system, unavailable)

[0323] @available(media operating system, unavailable)

[0324] @available(wearable operating system, unavailable)

[0325] @available(mixed reality operating system 1.0, *)

[0326] @propertyWrapper

[0327] public struct PhysicalMetric<Value> {

[0328]

[0329] /// A value in points in the coordinate system of the associated scene.

[0330] public var wrappedValue: Value { get }

[0331]

[0332] /// Creates a value that maps the specified single physical length measurement,

[0333] /// in the specified unit, to the corresponding value in points in the

[0334] /// associated scene.

[0335] public init(_ value: CGFloat, in unit: UnitLength) where Value == CGFloat

[0336]

[0337] /// Same doc comment as above, with minor corresponding edits, for:

[0338] public init(wrappedValue value: Value, from unit: UnitLength) where Value :
VectorArithmetic

[0339] public init(wrappedValue point: CGPoint, from unit: UnitLength) where Value
== CGPoint

[0340] public init(wrappedValue size: CGSize, from unit: UnitLength) where Value
== CGSize

[0341] public init(wrappedValue rect: CGRect, from unit: UnitLength) where Value
== CGRect

[0342] public init(wrappedValue point: Point3D, from unit: UnitLength) where Value
== Point3D

[0343] public init(wrappedValue size: Size3D, from unit: UnitLength) where Value ==
Size3D

[0344] public init(wrappedValue rect: Rect3D, from unit: UnitLength) where Value ==
Rect3D

[0345] }

[0346] For example:

[0347]

[0348] ```swift

[0349] struct PhysicalChessboardView: View {

[0350] @PhysicalMetric(from: .inches)

[0351] var singleChessboardSquareSide = 2.5

[0352]

[0353] @PhysicalMetric(from: .inches)

[0354] var fullChessboardAndWoodenBorderSize = CGSize(width: 22, height: 22)

[0355]

[0356] ...

[0357] }

[0358] ```

[0359] ```

[0360] In some examples, the values of these properties will be in points, and will match the given physical lengths provided. In some examples, if the points-per-meter value of a dynamicallyscaled scene changes, the view using this property wrapper will be invalidated. Using this environment key outside of a View that is in a view hierarchy is not supported; see below.

[0361] In some examples, both PhysicalMetric and PhysicalMetricsConverter require the context of a connected scene to obtain the correct points-per-meter ratio. These types are accessible in contexts that do not have that information. Namely: Some of these contexts, like a Scene or App struct, are inherently incorrect. In View s, this information is accessible when the content is being hosted in the view hierarchy. In some examples, these APIs are documented as not supported outside of a View . We leave what this means to be implementation-defined, but not all these situations are detectable or avoidable. In practice: If

we detect definite improper usage (e.g., outside a View), we reserve the right to produce a precondition failure on detection. If we detect usage that is proper, but in an improper context (e.g., a View not hosted in the view hierarchy), we reserve the right to alert the user by producing a fault in the Xcode runtime issues log. In these cases, we will use a default or hardcoded points-per-meter ratio. This signs us up for continually improving our logic so that we do not produce false positives.

[0362] These changes are additive and do not impact existing code. The `PhysicalMetricsConverter` type provides a basic operation; any further directions will depend on feedback, and will likely concentrate in the area of further improvements to scene sizing or additional conveniences.

[0363] All of the names above have several alternatives that can be explored in review. Support for initial sizing (`defaultSize`) with physical metrics for 2.5D windows has been deferred to a later release. We plan to use linked-on-or-after checks in those versions of the OS to not interfere with the current semantics, which do nothing to anything except for volumetric scenes. We could provide a raw points-Per-Meter value through the environment, as we do as a SPI. I believe this would make the use sites less readable, by requiring explicit math at each site with the possibility of type confusion between values that are meant to be in real-world lengths and values that are meant to be in points. In the current design, despite using the convenience of existing points-based types, every single real-world measurement is accompanied visually by an adjacent in: `UnitLength` parameter, so that it is clear at the call site that the particular value is not in points.

[0364] We could not provide any access to lengths at all. This would hamper several use cases; for example, it could cause developers to hardcode our current rules for point sizing in statically sized volumetric scenes, which we want to avoid. Dynamic interfaces to these values allow us to design different kinds of scene behaviors for future XR use cases without invalidating the assumptions of existing view code. `PhysicalMetric` is provided as a convenience; it can be replaced by advising use of `PhysicalMetricsConverter` instead. In a previous version of this draft, `PhysicalMetric` and `PhysicalMetricsConverter` were simply not available to use in types not associated with a scene's view contents. However, making the return types optional was ruled out as unergonomic. This may make code like the following legal to write and compile, but incorrect:

```

[0365]     ``swift
[0366]     struct MyVolumetricApp: App {
[0367]         @PhysicalMetric(from: .meters)
[0368]         var oneMeter = 1
[0369]
[0370]         var body: some Scene {
[0371]             WindowGroup {
[0372]                 MyVolumeView()
[0373]             }
[0374]             .windowStyle(.volumetric)
[0375]             .defaultSize(width: oneMeter, height: oneMeter, depth: oneMeter)
[0376]         }
[0377]     }
[0378]     ``

```

[0379] If this code worked, it might have seemed reasonable. If we are able to detect this case, we may fire a precondition failure; see above.

[0380] Just as a window scene creates space for a 2D application to exist, a volumetric scene creates space for a 3D application. One example use for a volumetric scene is for a future iteration of the Maps application. Most people typically use Maps in its 2D view, looking directly down onto the map. On existing platforms, they can then put Maps into a 3D view, where the camera pans down, and the map seems to come to life: detailed models of the buildings rise out of the ground, letting people explore cities in a new way. Mixed reality operating system, being an inherently 3D platform, is a natural fit to extend this functionality: this 3D map would render as if it were a table-sized model, in front of you. To get a different perspective on it, you simply move around it, letting you peek behind buildings and peer

down from above. On mixed reality operating system, 3D content can already be displayed within existing windows, but this content is not treated any differently from a typical 2.5D application window. One immediately obvious difference is in scale. Application windows, in order to keep text and UI legible, appear on the platform with a consistent angular size, meaning no matter how close or far away the window is, it will always occupy the same amount of your field of vision. A window could appear to be size of your computer monitor just in front of you, or the size of a billboard, several hundred feet away. This makes sense for most applications in order to preserve their usability, but is incongruous when attempting to represent things that more closely mirror real life. A 3D Chess game, for example, should not suddenly dwarf the table it's on when you move it across the room. As such, a volumetric scene confers both qualitative differences at the app level (this is a new class of app which is fundamentally built for 3D, rather than being adapted from the existing world of fundamentally 2D software), as well as quantitative differences in how the system presents these apps (applying a different scaling technique, a different set of visual effects, and a different style of chrome).

[0381] A developer specifies that a Window, WindowGroup, or DocumentGroup is Volumetric with a new window style:

[0382] ```swift

[0383] struct ChessApp: App {

[0384] var body: some Scene {

[0385] WindowGroup(id: "chessboard") {

[0386] View{ ... }

[0387] }

[0388] .windowStyle(.volumetric)

[0389] }

[0390] }

[0391] ```

[0392] By applying this window style to any existing window type, the backing UIWindowScene gets created in the system shell with new volumetric state and associated behaviors. Volumetric scenes in v1 support static sizing, composing with the defaultSize() modifier:

[0393]

[0394] ```swift

[0395] WindowGroup() { ... }

[0396] .windowStyle(.volumetric)

[0397] .defaultSize(width:400, height:400, depth:400)

[0398] ```

[0399] If the developer does not supply a size with this modifier, the system will apply a standard default size for a volumetric window, to be determined by HI. A volumetric scene can also be requested via the same openWindow environment property that developers are used to:

[0400] ```swift

[0401] struct WindowView: View {

[0402] @Environment(\.openWindow) private var openWindow

[0403]

[0404] var body: some View {

[0405] Button("Play Chess!") {

[0406] openWindow("chessboard")

[0407] }

[0408] }

[0409] }

[0410] ``

[0411] When a volumetric scene is opened, its placement is handled by the system shell. There is no specific attempt to place the scene on the floor or other horizontal surface, and the scene is infinitely repositionable by the user without any sort of snapping behavior to real-world objects.

[0412] While intended to host 3D content, usually created with three-dimensional kit and hosted in a View , most 2D user interface elements can still be used. Like the 3D content, the 2D UI will also draw in true scale, which could impact interactability depending on its distance from the user. This is a limitation for v1, and something we plan to address in the future; see "Mixed Scale" under Future Directions for details. Views laid out in the volumetric scene are centered on the X and Y axes by default, flowing vertically like in a 2D application. They occupy what is effectively the vertical plane of the scene's space, offset from the back of the volume by a fixed amount. These views can be repositioned arbitrarily using standard two-dimensional framework layout API, including the 3D layout extensions added for mixed reality operating system which are out of the scope of this proposal. Sheets and modal presentations require special consideration. In typical 2.5D applications, presenting a modal sheet displaces the main app window backward, with the sheet coming up in front. If multiple modals are presented, this appears as a stack, with the most recent modal on top. For volumetric scenes, this displacement behavior makes less sense: because the framework is not aware of the direction the user is relative to the scene (their "cardinality" relative to the scene), modals cannot be shown on the user-facing side, the way that chrome does. Displacement also makes less sense for volumetric content than it does for a floating platter window. Because of these limitations, the current plan is to not allow developers to present modal sheets in volumetric apps. An exception will be made for specific system UI such as Game Center and Apple Pay, where we can more precisely control its content and appearance.

[0413] Phone applications have their origin in the top-left corner of the app, with the X axis running horizontally to the right, and the Y axis vertically downward. This same convention is used for platter-style applications on mixed reality operating system. Volumetric applications extend this naturally by putting the origin at the top-left-back corner

of the box, relative to the user. The X axis projects to the user's right, the Y axis toward the floor, and the Z axis outward toward the user. Developers work in points when laying out content in applications. On 2D systems, a point is a pixel-agnostic unit of linear measure, with its origins in physical printing. Because our devices have high-DPI screens of varying pixel densities, working in points is a useful abstraction for expressing the size of onscreen content. The exact ratio of a point to a pixel varies based on display zoom and pixel density, but is generally around a 2:1 to 3:1 scale. N301 presents a new model where application content is entirely divorced from the pixel-representation on screen, instead representing some factor of real distance. For dynamic scale applications, a point is the distance of 1 millimeter at one meter away from the eye, representing a specific amount of angular space. The exact points-per-meter value is calculated during initial scene creation and placement, and then updated after each repositioning of the window. Developers should not need to concern themselves with this, letting them define content in terms of points as they are already used to, and knowing the content will be functional at any distance away. For true scale, a point is equal to 1 millimeter. See "Future Considerations" for discussion on accepting real-world measures.

[0414] The creation of all scenes on phoneOS-derived platforms is governed by the system shell. The core of the mechanical changes for volumetric scenes are all applied at the system shell level. During scene creation, the shell picks up the volumetric scene request and applies the following changes: Volumetric scenes get true scale, not dynamic scale, Volumetric scenes get user-facing chrome, which changes side based on the user's position, and Volumetric scenes do not fade out when viewed from angles other than the front.

[0415] Because UIKit's means of exposing this type is through a new scene session role, `UISceneSession.Role.windowApplicationVolumetric`, two-dimensional framework developers will also be able to see this value by using a `UIApplicationDelegateAdaptor`, allowing for unique scene customization from their scene delegate callbacks.

[0416] Add `VolumetricWindowStyle.swift`:

[0417] ````swift`

[0418] `/// A window style which creates a 3D volumetric window.`

[0419] `///`

```
[0420]    /// Use ``WindowStyle/volumetric`` to construct this style.

[0421]    @available(phoneOS, unavailable)

[0422]    @available(computerOS, unavailable)

[0423]    @available(wearableOS, unavailable)

[0424]    @available(mediaOS,unavailable)

[0425]    @available(mixed reality operating system 1.0, *)

[0426]    public struct VolumetricWindowStyle: WindowStyle {

[0427]        // Implementation is package internal

[0428]    }

[0429]

[0430]    @available(phoneOS, unavailable)

[0431]    @available(computerOS, unavailable)

[0432]    @available(wearableOS, unavailable)

[0433]    @available(mediaOS,unavailable)

[0434]    @available(mixed reality operating system 1.0, *)

[0435]    extension WindowStyle where Self == VolumetricWindowStyle {

[0436]        /// A window style which creates a 3D volumetric window.

[0437]        ///

[0438]        /// Volumetric windows are intended to host 3D content within a bounded

[0439]        /// region. Content in this region will get true scale by default, and allows

[0440]        /// the user to view the content from all angles without the scene contents
```

[0441] /// fading out. Volumetric scenes are intended for contained content and

[0442] /// should not be used for immersive environments.

[0443] public static var volumetric: VolumetricWindowStyle { get }

[0444] }

[0445] ```\n

[0446] Mark existing DefaultWindowStyle and .automatic option available. On mixed reality operating system it resolves to the typical 2.5D "platter" app style. Internally, this relates to the PlatterWindowStyle SPI. We are working with marketing to determine the final name for this type of app window, and exposing a concrete WindowStyle here will come in a separate addendum proposal once that name is finalized.

[0447] ```swift

[0448] @available(phoneOS, unavailable)

[0449] @available(computerOS 11.0, *)

[0450] @available(mediaOS, unavailable)

[0451] @available(wearableOS, unavailable)

[0452] @available(mixed reality operating system 1.0, *)

[0453] extension WindowStyle where Self == DefaultWindowStyle {

[0454] /// The default window style.

[0455] public static var automatic: DefaultWindowStyle { get }

[0456] }

[0457]

[0458] /// The default window style.

[0459] ///

[0460] /// You can also use ``WindowStyle/automatic`` to construct this style.

[0461] @available(phoneOS, unavailable)

[0462] @available(computerOS 11.0, *)

[0463] @available(mediaOS, unavailable)

[0464] @available(wearableOS, unavailable)

[0465] @available(mixed reality operating system 1.0, *)

[0466] public struct DefaultWindowStyle: WindowStyle {

[0467]

[0468] public init() {...}

[0469] }

[0470] ``` Mark WindowStyle as available on mixed reality operating system:

[0471] ```swift

[0472] @available(phoneOS, unavailable)

[0473] @available(computerOS 11.0, *)

[0474] @available(mediaOS, unavailable)

[0475] @available(wearableOS, unavailable)

[0476] @available(mixed reality operating system 1.0, *)

[0477] public protocol WindowStyle {...}

[0478]

[0479] @available(phoneOS, unavailable)

[0480] @available(computerOS 11.0, *)

```
[0481]    @available(mediaOS, unavailable)

[0482]    @available(wearableOS, unavailable)

[0483]    @available(mixed reality operating system 1.0, *)

[0484]    extension Scene {

[0485]        /// Sets the style for windows created by this scene.

[0486]        public func windowStyle<S: WindowStyle>(_ style: S)

[0487]            -> some Scene {...}

[0488]    }

[0489]    ``Volumetric windows support static sizing with the defaultSize modifier:

[0490]    ``swift

[0491]    @available(phoneOS, unavailable)

[0492]    @available(computerOS, unavailable)

[0493]    @available(wearableOS, unavailable)

[0494]    @available(mediaOS,unavailable)

[0495]    @available(mixed reality operating system 1.0, *)

[0496]    extension Scene {

[0497]        /// Sets a default size for a volumetric window.

[0498]        ///

[0499]        /// Use this modifier to indicate the default initial size for a new 3D

[0500]        /// window created from a ``Scene`` using ``VolumetricWindowStyle``:

[0501]        ///
```

```
[0502]    /// WindowGroup {
[0503]    ///     ContentView()
[0504]    /// }
[0505]    /// .windowStyle(.volumetric)
[0506]    /// .defaultSize(width: 600, height: 400, depth: 600)
[0507]    ///
[0508]    /// Each parameter is specified in points, which translates to 1 millimeter
[0509]    /// for volumetric scenes at standard system scale. The size of a volumetric
[0510]    /// scene is immutable after creation.
[0511]    ///
[0512]    /// DefaultSize is a no-op for other window styles on mixed reality operating
    system.
[0513]    ///
[0514]    /// - Parameter width: The default width for the created window.
[0515]    /// - Parameter height: The default height for the created window.
[0516]    /// - Parameter depth: The default depth for the created volumetric window.
[0517]    /// Depth is disregarded on scenes that don't support specifying depth.
[0518]    ///
[0519]    /// - Returns: A scene that uses a default size for new windows.
[0520]    public func defaultSize(
[0521]        width: CGFloat, height: CGFloat, depth: CGFloat
[0522]    ) -> some Scene {...}
```

[0523]

[0524] /// Sets a default size for a volumetric window.

[0525] ///

[0526] /// Use this modifier to indicate the default initial size for a new 3D

[0527] /// window created from a ``Scene`` using ``VolumetricWindowStyle``:

[0528] ///

[0529] /// WindowGroup {

[0530] /// ContentView()

[0531] /// }

[0532] /// .windowStyle(.volumetric)

[0533] /// .defaultSize(Size3D(width: 600, height: 400, depth: 600))

[0534] ///

[0535] /// Each parameter is specified in points, which translates to 1 millimeter

[0536] /// for volumetric scenes at standard system scale. The size of a volumetric

[0537] /// scene is immutable after creation.

[0538] ///

[0539] /// DefaultSize is a no-op for other window styles on mixed reality operating
system.

[0540] ///

[0541] /// - Parameter size: The default 3D size for the created window.

[0542] /// Depth is disregarded on scenes that don't support specifying depth.

[0543] ///

[0544] /// - Returns: A scene that uses a default size for new windows.

[0545] public func defaultSize(_ size: Size3D) -> some Scene { ... }

[0546] }

[0547] ```\n

[0548] In some examples, after creation, volumetric windows cannot be resized, and within any restrictions applied by the shell, the size specified by the developer in defaultSize will be the size given for the scene. Because the system shell has the right to not honor the given size, developers will need some mechanism to determine their size at runtime. We have added GeometryReader3D as SPI, and moving this to API is likely the supported path for determining the effective size of a volumetric window. Because volumetric windows cannot be resized after creation, the windowResizability() modifier is disregarded and acts as a no-op.

[0549] While the volumetric window style lays the groundwork for this class of application, there are additional behaviors which have been moved out to future releases: Mixed Scale: Developers will want to be able to present dynamic scale UI within a true scale volumetric application. For example, an info sheet popping out of a point of interest on a map, where the sheet is dynamic scale but the map remains true scale. This will require significant engineering effort and is out of scope for v1. The API for mixed scale could look like:

[0550] ```swift

[0551] // Marking a view as dynamic scale

[0552] View()

[0553] .scaling(.dynamic)

[0554]

[0555] // Specifying the scaling point

[0556] // This view will scale from its bottom-center

[0557] View()

[0558] .scaling(.dynamic, anchor: UnitPoint(x: 0.5, y: 1.0)

[0559] ``

[0560] Toolbar: Initial HI mock-ups for the feature included a floating pill-shaped toolbar below the content for UI elements. In some examples, this is built using the existing .toolbar() modifier in two-dimensional framework, but this is out of scope for v1. 3D Layout API: For v1, the 3D content displayed within a volumetric window will generally place three-dimensional kit content near the root of the app, presented within a View. In the future, we would like to offer more robust native 3D layout API to let developers build volumetric content with a toolkit more closely resembling a two-dimensional framework. User Resizability: In some examples, initially, volumetric windows will support fixed sizing using the defaultSize modifier. In some examples, support for resizing is supported. Presentations: As mentioned above, developer presentations will be mitigated and logged in v1. Longer-term, we want to allow presentations, but this will require additional engineering effort as well as HI work to determine how they should appear. In some examples, when volumetric becomes an available style on other platforms, presentations will work on them as well. Size in real units: In v1, the .defaultSize modifier takes points as the unit of measure. Particularly for volumetric scenes, since they run in true scale, developers may want to specify sizes in real-world units, rather than in points. To expose this to users, a new .defaultSize overload like the following is added:

[0561] ``swift

[0562] extension Scene {

[0563] public func defaultSize(

[0564] width: CGFloat, height: CGFloat, depth: CGFloat, in unit: UnitLength

[0565]) -> some Scene {}

[0566] }

[0567] ``

[0568] Expose a Concrete "Platter Window" Style: The .automatic style should always resolve to a concrete style when possible. For mixed reality operating system, this default style would create a typical dynamic-scale, plattered app. Marketing is currently trying to come up with a good term for this style of window, and the API should follow that name if possible. For the moment exposing this type via the .automatic style is fine. Brainstormed names for this style include: Pane, Platter, Flat, TwoDee, and Windowed.

[0569] There is no certain relation between wanting the volumetric scene behavior, and including a View. Because we want to add native 3D layout API in the future, there's no guarantee that the root view of a volumetric window will always be a View . And conversely, developers can also use View to present three-dimensional kit content within an otherwise traditional 2.5D application on the platform. This was rejected because it lacks composability. A volume app could be document-based, be a single window, or be a duplicatable window showing copies of the same contents. All these behaviors are fundamentally the same as the existing two-dimensional framework Scene types, WindowGroup, Window, and DocumentGroup and so it makes more sense to modify those types with a consistent modifier, rather than create a new set of parallel scene types. We want to avoid introducing new, special case API where we can. While this would be feasible, making windowStyle() available on mixed reality operating system gives us more flexibility for defining additional styles in the future. In some example, volumetric window style will be available on mixed reality operating system. It is worth some additional thought to how it could be used on other platforms as well. On computer operating system, a .volumetric window style could perhaps be used perhaps to configure a window with a 3D context for rendering. With WindowStyle being available on computer operating system and mixed reality operating system, it is worth considering if the modifier ought to be available everywhere, in the name of reducing conditional code for developers in cross-platform apps. However, because the specific styles available will still differ per-platform, this availability consideration should be considered orthogonally to exposing .volumetric on mixed reality operating system . mixed reality operating system is adding new support on UIWindowScene to display volumetric 3D content. In order to get a window scene that is configured for this, developers need a way to signal to the system shell that a given scene should gain these features. Developers also need to be able to determine at runtime whether a given scene is volumetric or not, for example in their scene:willConnectToSession: method, or when they receive the scene from another process. Volumetric scenes are built on top of the existing

UIWindowScene and do not themselves require additional UIKit API or a UIWindowScene subclass, as the majority of what defines a volumetric scene is applied by the system shell. For more details on what a volumetric scene entails, see two-dimensional framework - 105201522 Volumetric Scenes.

[0570] Add a new UISceneSessionRole called UIWindowSceneSessionRoleVolumetricApplication. In Swift, it is exposed as UISceneSession.Role.windowApplicationVolumetric. A scene session role is the perfect choice to indicate this kind of behavioral change. The session role is a semantic descriptor of the purpose of a scene, and does not require a UIScene subclass or additional overhead; it is a clear indicator the developer can use to signal to the system shell that the scene should be treated differently; and it is easily referenced at runtime off the scene session.

[0571] This role can be used to request a volumetric scene using the new UISceneSessionActivationRequest object (rdar://102112617 Scene Request Object API):
Objective-C: UISceneSessionActivationRequest *activationRequest =
[UISceneSessionActivationRequest [UIApplication.sharedApplication
activateSceneSessionForRequest:activationRequest].

[0572] ```swift

[0573] let request = UISceneSessionActivationRequest(role:
.windowApplicationVolumetric)

[0574] UIApplication.shared.activateSceneSession(for: request) { err in ... }

[0575] ```

[0576] It can also be used with an existing scene to determine whether it is a volumetric scene:

[0577] ```objective-c

[0578] if (windowScene.session.role ==
UIWindowSceneSessionRoleVolumetricApplication) {

[0579] // The scene is volumetric, do 3D things with it

[0580] ...

[0581] } else {

[0582] ...

[0583] }

[0584] ```

[0585] For two-dimensional framework users, this is also useful for those using a UIApplicationDelegateAdaptor.

[0586] Add the following to <UIKit/UIWindowScene.h>

[0587] ```swift

[0588] extension UISceneSession.Role {

[0589] @available(mixed reality operating system 1.0, *)

[0590] @available(phoneOS, unavailable)

[0591] @available(mediaOS, unavailable)

[0592] @available(computerOS, unavailable)

[0593] @available(wearableOS, unavailable)

[0594] static let windowApplicationVolumetric: UISceneSession.Role

[0595] }

[0596] ```

[0597] ```objective-c

[0598] /// A session role for volumetric 3D scenes.

[0599] UIKIT_EXTERN UISceneSessionRole const

UIWindowSceneSessionRoleVolumetricApplication API_AVAILABLE(mixed reality

operating system (1.0)) API_UNAVAILABLE(phoneOS, mediaOS, computerOS, wearableOS);

[0600] ```

[0601] In some examples, this change is additive, to enable new behaviors. Applications written for mixed reality operating system will need to ensure that anywhere they check session role, that they properly handle this session role in addition to the existing ones. In most cases, I would expect checks for UIWindowSceneSessionRoleApplication to want to also include UIWindowSceneSessionRoleVolumetricApplication, so mixed reality operating system apps should audit reused code that checks the session role and decide what the right course is for themselves.

[0602] Adding a new scene session role is a clear fit for the need to indicate a new type of scene with different behaviors. The alternative would be not providing an alternate public role, and having volumetric scenes also indicate UIWindowSceneSessionRoleApplication . This would be unideal, as we have clients requesting the ability to determine whether a scene they receive from out-of-process is a volumetric scene or "regular" window scene. The need to determine the type of scene at runtime is high enough that it warrants the creation of a new session role. Additionally, for developers who specify their scenes using the Scene Manifest in their Info.plist, they will use the volumetric session role to indicate volumetric scenes to the shell.

[0603] Originally, the role name included the term "Bounded" as a way to disambiguate from Boundless scenes which are also being introduced for Wolf.

[0604] Mixed reality operating system makes WindowResizability and its modifier available on the mixed reality operating system. The default value, if not specified (or if automatic is used, would be contentMinSize.

[0605] ```diff

[0606] @available(phoneOS 17.0, *)

[0607] @available(computerOS 13.0, *)

[0608] @available(mediaOS, unavailable)

[0609] @available(wearableOS, unavailable)

[0610] -@available(mixed reality operating system, unavailable)

[0611] +@available(mixed reality operating system 1.0, *)

[0612] extension Scene {

[0613] public func windowResizability(

[0614] _resizability: WindowResizability

[0615]) -> some Scene

[0616] }

[0617]

[0618] @available(phoneOS 17.0, *)

[0619] @available(computerOS 13.0, *)

[0620] @available(mediaOS, unavailable)

[0621] @available(wearableOS, unavailable)

[0622] -@available(mixed reality operating system, unavailable)

[0623] +@available(mixed reality operating system 1.0, *)

[0624] public struct WindowResizability: Sendable {

[0625] public static let automatic: Self

[0626] public static let contentSize: Self

[0627] public static let contentMinSize: Self

[0628] }

[0629] ````

[0630] Window Scenes on mixed reality operating system can be resized by the user or by the application (much like computerOS). A developer needs to be able to request a new size for their window scene. Add a platform specific geometry preference object UIWindowSceneGeometryPreferences. This API will start with being able to set the size (as a CGSize), minimumSize, maximumSize, and resizingRestriction (see below). The real values for these will be reflected in the effectiveGeometry object as they change.

[0631] The geometry preference can be used to specify scene size as well as restrictions for how the user may resize the scene.

[0632] ```objective-c

[0633] UIWindowSceneGeometryPreferences *pref =
[[UIWindowSceneGeometryPreferencesReality alloc] initWithSize:CGSizeMake(500, 500)];

[0634]

[0635] // Sets the minimumSize to 400, 400

[0636] pref.minimumSize = CGSizeMake(400, 400)

[0637]

[0638] // Enforce that user resizes to the current programatic set aspect ratio (in this case 1:1)

[0639] pref.resizingRestrictions = UIWindowSceneResizingRestrictionsUniform

[0640]

[0641] [windowScene requestGeometryUpdateWithPreferences:nil
errorHandler:^(NSError * _Nonnull error) {

[0642] // handle error

[0643] }];

[0644] ```

[0645] ```swift

```
[0646]     scene.requestGeometryUpdate(.Reality(size: .init(width: 500, 500),
minimumSize: CGMakeSize(400, 400), resizingRestrictions: .uniform)) { err in

[0647]         // handle error

[0648]     }

[0649]     ``

[0650]     Add the following to <UIKit/UIWindowSceneGeometry.h>

[0651]     ``objective-c

[0652]     typedef NS_ENUM(NSInteger, UIWindowSceneResizingRestrictions) {

[0653]         /// App has no preference on user resize

[0654]         UIWindowSceneResizingRestrictionsUnspecified = 0,

[0655]

[0656]         /// The user cannot resize the scene

[0657]         UIWindowSceneResizingRestrictionsNone,

[0658]

[0659]         /// User resizes are restricted to the current aspect ratio

[0660]         UIWindowSceneResizingRestrictionsUniform,

[0661]

[0662]         /// In some examples, user resizes are restricted by the system and other
restrictions put in place

[0663]         UIWindowSceneResizingRestrictionsFreeform

[0664]     } API_AVAILABLE(mixed reality operating system (1.0));

[0665]
```

[0666] @interface UIWindowSceneGeometry ()

[0667]

[0668] /// The current app specified minimumSize. A value of 0,0 is returned if a minimum is not set by the application

[0669] @property (nonatomic, readonly) CGSize minimumSize
API_AVAILABLE(mixed reality operating system (1.0));

[0670]

[0671] /// The current app specified maximumSize. A value of
CGFLOAT_MAX,CGFLOAT_MAX is returned if a maximum is not set by the application

[0672] @property (nonatomic, readonly) CGSize maximumSize
API_AVAILABLE(mixed reality operating system (1.0));

[0673]

[0674] /// The current app specified resizingRestriction. Default value
UIWindowSceneResizingRestrictionsUnspecified

[0675] @property (nonatomic, readonly) UIWindowSceneResizingRestrictions
resizingRestrictions API_AVAILABLE(mixed reality operating system (1.0));

[0676]

[0677] @end

[0678]

[0679] ```

[0680] ```objective-c

[0681] typedef NS_ENUM(NSUInteger, UIWindowSceneResizingRestrictions) {

[0682] /// App has no preference on user resize

```
[0683]    UIWindowSceneResizingRestrictionsUnspecified = 0,

[0684]

[0685]    /// The user cannot resize the scene

[0686]    UIWindowSceneResizingRestrictionsNone,

[0687]

[0688]    /// User resizes are restricted to the current aspect ratio

[0689]    UIWindowSceneResizingRestrictionsUniform,

[0690]

[0691]    /// In some examples, user resizes are restricted by the system and other
restrictions put in place

[0692]    UIWindowSceneResizingRestrictionsFreeform

[0693] } API_AVAILABLE(mixed reality operating system (1.0));

[0694]

[0695] @interface UIWindowSceneGeometry ()

[0696]

[0697]    /// The current app specified minimumSize. A value of 0,0 is returned if a
minimum is not set by the application

[0698]    @property (nonatomic, readonly) CGSize minimumSize
API_AVAILABLE(mixed reality operating system (1.0));

[0699]

[0700]    /// The current app specified maximumSize. A value of
CGFLOAT_MAX,CGFLOAT_MAX is returned if a maximum is not set by the application
```

[0701] @property (nonatomic, readonly) CGSize maximumSize
API_AVAILABLE(mixed reality operating system (1.0));

[0702]

[0703] /// The current app specified resizingRestriction. Default value
UIWindowSceneResizingRestrictionsUnspecified

[0704] @property (nonatomic, readonly) UIWindowSceneResizingRestrictions
resizingRestrictions API_AVAILABLE(mixed reality operating system (1.0));

[0705]

[0706] @end

[0707]

[0708] ``

[0709] Create a new public header
<UIKit/UIWindowSceneGeometryPreferencesReality.h>

[0710] ``objective-c

[0711] /// Used as the value for a dimension of a size related preference when wanting to
leave it unchanged.

[0712] UIKIT_EXTERN const CGFloat UIProposedSceneSizeNoPreference
API_AVAILABLE(mixed reality operating system (1));

[0713]

[0714] UIKIT_EXTERN API_AVAILABLE(mixed reality operating system (1.0))

[0715] @interface UIWindowSceneGeometryPreferencesReality :
UIWindowSceneGeometryPreferences

[0716]

[0717] - (instancetype)init NS_DESIGNATED_INITIALIZER;

[0718]

[0719] /// Creates a geometry preference with a specific size (specifying
UIProposedSceneSizeNoPreference for any dimension of size, will specify no preference,
keeping that dimension the same if possible)

[0720] - (instancetype)initWithSize:(CGSize)size;

[0721]

[0722] /// The preferred size, minimumSize, maximumSize, and resizingRestrictions. Any
dimension set to UIProposedSceneSizeNoPreference will retain the existing values.

[0723] - (instancetype)initWithSize:(CGSize)size minimumSize:(CGSize)minimumSize
maximumSize:(CGSize)maximumSize
resizingRestrictions:(UIWindowSceneResizingRestrictions)resizingRestrictions;

[0724]

[0725] /// Create a new preference using the existing effective geometry

[0726] - (instancetype)initWithCurrentEffectiveGeometry:(UIWindowSceneGeometry
*)effectiveGeometry;

[0727]

[0728] /// The preferred system size. Use UIProposedSceneSizeNoPreference to use
existing value

[0729] @property (nonatomic, assign) CGSize size NS_REFINED_FOR_SWIFT;

[0730]

[0731] /// The preferred minimum size. UIProposedSceneSizeNoPreference to use default
value

[0732] @property (nonatomic, assign) CGSize minimumSize
NS_REFINED_FOR_SWIFT;

[0733]

[0734] /// The preferred maximum size. UIProposedSceneSizeNoPreference for a given
to use default value

[0735] @property (nonatomic, assign) CGSize maximumSize
NS_REFINED_FOR_SWIFT;

[0736]

[0737] /// The user resizable restrictions on the window scene

[0738] @property (nonatomic, assign) UIWindowSceneResizingRestrictions
resizingRestrictions NS_REFINED_FOR_SWIFT;

[0739] @end

[0740] ``` `swift

[0741] @available(mixed reality operating system 1.0, *)

[0742] class UIWindowScene.GeometryPreferences.Reality:
UIWindowScene.GeometryPreferences {

[0743] var size: CGSize?

[0744] var minimumSize: CGSize?

[0745] var maximumSize: CGSize?

[0746] var resizeRestrictions: ResizingRestrictions?

[0747]

[0748] init(effectiveGeometry: UIWindowScene.Geometry)

[0749] init(size: CGSize? = nil, minimumSize: CGSize? = nil, maximumSize:
CGSize? = nil, resizeRestrictions: ResizingRestrictions? = nil)

[0750] }

[0751] }

[0752] '''

[0753] In some examples, this API adds platform specific preferences
 UIWindowSceneGeometryPreferencesReality does not have impact on other platforms.
 Presently, the additions to UIWindowSceneGeometry are mixed reality operating system for .
 Catalyst has systemFrame which is different enough that we don't think makes sense for
 mixed reality operating system.. There will be other API proposals to deprecate
 UISceneSizeRestrictions and move that functionality over to their platforms respective
 UIWindowSceneGeometryPreferences . rdar://108554504 (Deprecate
 UISceneSizeRestrictions and move content over to the platform specific
 UIWindowSceneGeometryPreferences) With N301 and mixed reality operating system, we
 are bringing two-dimensional framework into the three dimensional world, enabling new
 kinds of applications and experiences built using 3D content, gestures, and effects. A
 common use case for working in a 3D environment is to apply true 3D transforms to content.
 On mixed reality operating system, this is often used to: Add Z offsets to views to provide a
 sense of visual hierarchy and emphasis on select elements, even in a 2.5D windowed
 application Apply 3D rotations, scales, and offsets to content, such as when manipulating a
 3D model with a gesture because two-dimensional framework on mixed reality operating
 system supports use of 3D content through APIs like Model3D and View (proposed
 separately), as well as 3D scenes like volumetric windows and unbounded Stages, it is also
 important to be able to apply these 3D transforms relative to a view's 3D size, such as to
 rotate a 3D model about its center. In some examples, , mixed reality operating system
 supports applying 3D transforms both to 3D content (via three-dimensional kit) and to 2D
 views. When a 3D transform is applied to a 2D view, the view is rendered in a separate mesh
 and texture through a process known as "separation" baked into Core Animation and three-
 dimensional kit, which is leveraged internally to two-dimensional framework . We aim to
 expose these capabilities publicly in two-dimensional framework on mixed reality operating
 system with a mix of new APIs and new behaviors on existing APIs.

[0754] A suite of new View and VisualEffect modifiers for applying 3D transform effects
 to a view A new .visualEffect3D modifier, analogous to .visualEffect while providing access
 to a view's 3D geometry Amending the existing rotation3DEffect modifiers to provide a
 clearer story differentiating perspective rotations from true 3D rotations A new UnitPoint3D
 type, analogous to UnitPoint , used throughout the above APIs.

[0755] First, modifiers perform true 3D transform effects like scales, rotations, and translations on two-dimensional framework views. These are exposed both as View modifiers and as VisualEffect modifiers, as with existing transform effect modifiers like `.scaleEffect()`. The simplest and most commonly used 3D transform effect is to apply a Z offset to a view. These are used throughout mixed reality operating system, such as to present sheets visually in front of its presenting view, or to have controls lift up to meet the user's finger on direct touch hover. For example:

[0756] ````swift`

[0757] `Capsule()`

[0758] `.offset(z: hovering ? 10 : 0)`

[0759] We can also apply true 3D rotations using new variants of the `.rotation3DEffect()` modifier, which is particularly useful for orienting 3D content:

[0760]

[0761] ````swift`

[0762] `struct Earth: View {`

[0763] `@GestureState var rotation: Rotation3D = .identity`

[0764] `var body: some View {`

[0765] `Model3D("earth")`

[0766] `.rotation3DEffect(rotation)`

[0767] `.gesture(RotateGesture3D().updating($rotation) {`

[0768] `...`

[0769] `})`

[0770] `}`

[0771] `}`

[0772] Note: The Rotation3D type is defined in the Spatial framework, as are other 3D geometry APIs used throughout this proposal like AffineTransform3D and Size3D. The existing `.scaleEffect()` modifier also works automatically with 3D content when using a uniform scale factor, applying that same scale in the Z axis. This enables intuitive behavior in the common case, like with applying a magnify gesture:

```
[0773]     ``swift
[0774]     struct Earth: View {
[0775]         @GestureState var scale = 1.0
[0776]         var body: some View {
[0777]             Model3D("earth")
[0778]                 .scaleEffect(scale)
[0779]                 .gesture(MagnifyGesture().updating($scale) {
[0780]                     ...
[0781]                 })
[0782]         }
[0783]     }
[0784]     ``
```

[0785] If more granular control is needed, the new `.scaleEffect(x:y:z:anchor:)` modifier can apply non-uniform 3D scale effects as well. Finally, modifiers are added to apply a custom AffineTransform3D 4x3 transform to a view via the `.transform3DEffect()` modifier, analogous to the `.transformEffect()` modifier for applying a 3x3 ProjectionTransform. Because these modifiers are also available as VisualEffects, they can be used within APIs like the scrollTransition modifier — for example, to apply 3D effects to views in a ScrollView:

```
[0786]     ``swift
```

```

[0787] struct ContentView: View {
[0788]     var body: some View {
[0789]         ScrollView {
[0790]             ForEach(City.all) { city in
[0791]                 CityCard(city)
[0792]                 .scrollTransition { content, phase in
[0793]                     content
[0794]                     .offset(z: phase.isIdentity ? -10 : 0)
[0795]                 }
[0796]             }
[0797]         }
[0798]     }
[0799] }
[0800] ```

```

[0801] In some examples, to create custom 3D geometry effects, it is important to be able to read the 3D geometry of the view, such as to apply a transform about a 3D anchor point. The new `.visualEffect` modifier provides access to a `GeometryProxy` within its closure. In some examples, this includes the view's 2D size. To address this, introduce a `.visualEffect3D` modifier, which behaves like `.visualEffect` but instead provides a `GeometryProxy3D` (as proposed here). This enables more advanced effects, like a recessing effect that is proportional to the depth of a `Model3D`:

```

[0802] ```swift
[0803] struct RecedingModelExample: View {

```

```

[0804]     @State var receded = false
[0805]     var body: some View {
[0806]         Model3D("starship")
[0807]             .visualEffect3D { geometry, content in
[0808]                 content
[0809]                     .offset(z: receded ? -geometry.depth : 0)
[0810]                     .scaleEffect(recessed ? 0.5 : 1.0)
[0811]             }
[0812]         Toggle("Receded", isOn: $receded)
[0813]     }
[0814] }
[0815] ```

```

[0816] As noted earlier, using the rotation3DEffect name to generally represent a true 3D rotation on mixed reality operating system is advantageous, as this most accurately reflects the behavior expressed by the modifier name. To model the behavior of a true 3D rotation, new overloads with 3D types like Rotation3D are introduced:

```

[0817] ```swift
[0818] extension View {
[0819]     public func rotation3DEffect(
[0820]         _ angle: Angle, axis: RotationAxis3D, anchor: UnitPoint3D = .center)
[0821]     -> some View
[0822]     public func rotation3Deffect(

```

```

[0823]         _ rotation: Rotation3D, anchor: UnitPoint3D = .center)
[0824]     -> some View
[0825]     public func rotation3Deffect(
[0826]         _ angle: Angle, axis: (x: CGFloat, y: CGFloat, z: CGFloat),
[0827]         anchor: UnitPoint3D = .center
[0828]     ) -> some View
[0829] }
[0830] ```

```

[0831] Note that unlike the existing rotation3DEffect modifiers, in some examples, these overloads do not take a perspective value, as the concept does not have a sensible analog with a true 3D rotation. However, there are valid use cases on mixed reality operating system for preserving the existing behavior of rotation3DEffect on other platforms, where a perspective rotation is applied to the view as a 2D projection, without creating true 3D geometry. To support these use cases, a new perspectiveRotationEffect modifier is added:

```

[0832]
[0833] ```swift
[0834] extension View {
[0835]     public func perspectiveRotationEffect(_ angle: Angle,
[0836]         axis: (x: CGFloat, y: CGFloat, z: CGFloat),
[0837]         anchor: UnitPoint3D = .back, perspective: CGFloat = 1)
[0838]     -> some View
[0839] }
[0840] ```

```

[0841] To help further disambiguate between perspective rotations and true 3D rotations, hard-deprecating the existing rotation3DEffect modifiers on mixed reality operating system and directing developers to use perspectiveRotationEffect instead is proposed. In some examples, this affects apps that specify a perspective value or anchor explicitly. In some examples, when an angle and axis are specified, the new rotation3DEffect overload is preferred instead when compiling for the mixed reality operating system SDK and automatically results in applying a true 3D rotation:

[0842] ```swift

[0843] SomeView()

[0844] .rotation3DEffect(.degrees(45), axis: (x: 1, y: 0, z: 0))

[0845] ```

[0846] Note that this approach does not propose any changes to the 3D rotation APIs for other platforms at this time. See the Future directions section for additional thoughts on that subject.

[0847] To support the above APIs, a new UnitPoint3D type, serving as a 3D analog to the existing UnitPoint type is introduced. UnitPoint3D provides static members for common points, with additional variants to express the front, back, and center of a given 2D point. For example, UnitPoint3D.topLeadingBack represents the point at the back of the view in Z and the top-leading corner in X and Y, while UnitPoint3D.topLeading instead represents the center of the view in Z. Because there is some naming overlap with UnitPoint 's static members, some care must be taken with modifiers that can in some places take a UnitPoint3D and in others a UnitPoint . For instance, the new .scaleEffect(x:y:z:anchor:) modifier takes a UnitPoint3D , and is marked as @_disfavoredOverload to ensure that writing code like this:

[0848] ```swift

[0849] Color.red.scaleEffect(x: 2, y: 1, anchor: .topLeading)

[0850] ```

[0851] Biases towards using the existing modifier that takes a UnitPoint and no z parameter, preserving existing semantics and expectations for apps that are also built for other platforms.

[0852] ```swift

[0853] @available(phoneOS, unavailable)

[0854] @available(computerOS, unavailable)

[0855] @available(wearableOS, unavailable)

[0856] @available(mediaOS, unavailable)

[0857] @available(mixed reality operating system 1.0, *)

[0858] extension View {

[0859] /// Brings a view forward in Z by the provided distance in points.

[0860] ///

[0861] /// - Parameters:

[0862] /// - distance: The distance to extrude the view forward in Z, in points.

[0863] /// - Returns: A view that is extruded forward in Z by `distance`.

[0864] @_alwaysEmitIntoClient

[0865] public func offset(z: CGFloat) -> some View

[0866]

[0867] /// Rotates this view by the specified angle about an axis and anchor point.

[0868] ///

[0869] /// The original dimensions of the view are considered to be unchanged by

[0870] /// rotating the contents.

```
[0871]    ///
[0872]    /// For example, to rotate a 3D model by 45° about the Y axis:
[0873]    ///
[0874]    /// ```
[0875]    /// Model3D("robot")
[0876]    ///     .rotationEffect(.degrees(45), axis: .y)
[0877]    /// ```
[0878]    ///
[0879]    /// - Parameters:
[0880]    ///   - angle: The angle by which to rotate this view.
[0881]    ///   - axis: The axis about which to rotate this view.
[0882]    ///   - anchor: The anchor point about which to rotate this view. Defaults to
[0882]    center.
[0883]    /// - Returns: A view that rotates this view by `angle` about `axis` and `anchor`.
[0884]    @_alwaysEmitIntoClient
[0885]    public func rotation3DEffect(
[0886]        _ angle: Angle, axis: RotationAxis3D, anchor: UnitPoint3D = .center)
[0887]    -> some View
[0888]
[0889]    /// Rotates this view by a specific amount.
[0890]    ///
[0891]    /// The original dimensions of the view are considered to be unchanged by
```

```
[0892]    /// rotating the contents.

[0893]    ///

[0894]    /// For example, to rotate a 3D model by 45° about the Y axis:

[0895]    ///

[0896]    /// ```

[0897]    /// Model3D("robot")

[0898]    ///     .rotationEffect(Rotation3D(angle: .degrees(45), axis: .y))

[0899]    /// ```

[0900]    ///

[0901]    /// - Parameters:

[0902]    ///   - rotation: The rotation by which to rotate the view.

[0903]    ///   - anchor: The anchor point about which to rotate this view. Defaults to
center.

[0904]    /// - Returns: A view that rotates this view by `rotation` about `anchor`.

[0905]    @_alwaysEmitIntoClient

[0906]    public func rotation3DEffect(

[0907]        _ rotation: Rotation3D, anchor: UnitPoint3D = .center)

[0908]    -> some View

[0909]

[0910]    /// Rotates this view by the specified angle about an axis and anchor point.

[0911]    ///

[0912]    /// The original dimensions of the view are considered to be unchanged by
```

```
[0913]    /// rotating the contents.

[0914]    ///

[0915]    /// For example, to rotate a 3D model by 45° about the Y axis:

[0916]    ///

[0917]    /// ```

[0918]    /// Model3D("robot")

[0919]    ///     .rotation3DEffect(.degrees(45), axis: .y)

[0920]    /// ```

[0921]    ///

[0922]    /// - Parameters:

[0923]    ///   - angle: The angle by which to rotate this view.

[0924]    ///   - axis: The axis about which to rotate this view.

[0925]    ///   - anchor: The anchor point about which to rotate this view. Defaults to
center.

[0926]    /// - Returns: A view that rotates this view by `angle` about `axis` and `anchor`.

[0927]    @_alwaysEmitIntoClient

[0928]    public func rotation3DEffect(

[0929]        _ angle: Angle, axis: (x: CGFloat, y: CGFloat, z: CGFloat),

[0930]        anchor: UnitPoint3D = .center

[0931]    ) -> some View

[0932]

[0933]    /// Rotates this view's rendered output in three dimensions around the given
```

[0934] /// axis of rotation.

[0935] ///

[0936] /// Use `perspectiveRotationEffect(\_ :axis:anchor:anchorZ:perspective:)` to

[0937] /// rotate the view in three dimensions around the given axis of rotation,

[0938] /// and optionally, position the view at a custom display order and

[0939] /// perspective.

[0940] ///

[0941] /// In the example below, the text is rotated 45° about the `y` axis,

[0942] /// front-most (the default `zIndex`) and default `perspective` (`1`):

[0943] ///

[0944] /// Text("Rotation by passing an angle in degrees")

[0945] /// .perspectiveRotationEffect(.degrees(45), axis: (x: 0, y: 1, z: 0))

[0946] /// .border(Color.gray)

[0947] ///

[0948] /// ![A screenshot showing the rotation of text 45 degrees about the

[0949] /// y-axis.](two-dimensional framework -View-rotation3DEffect.png)

[0950] ///

[0951] /// - Parameters:

[0952] /// - angle: The angle at which to rotate the view.

[0953] /// - axis: The `x`, `y` and `z` elements that specify the axis of

[0954] /// rotation.

[0955] /// - anchor: The location with a default of ``UnitPoint/center`` that

[0956] /// defines a point in 3D space about which the rotation is anchored.

[0957] /// - anchorZ: The location with a default of `0` that defines a point in

[0958] /// 3D space about which the rotation is anchored.

[0959] /// - perspective: The relative vanishing point with a default of `1` for

[0960] /// this rotation.

[0961] @_alwaysEmitIntoClient

[0962] public func perspectiveRotationEffect(_ angle: Angle,

[0963] axis: (x: CGFloat, y: CGFloat, z: CGFloat),

[0964] anchor: UnitPoint3D = .back, perspective: CGFloat = 1)

[0965] -> some View

[0966]

[0967] /// Scales this view uniformly by the specified size in each dimension.

[0968] ///

[0969] /// The original dimensions of the view are considered to be unchanged by

[0970] /// scaling the contents. To change the dimensions of the view, use a

[0971] /// modifier like `frame()` instead.

[0972] ///

[0973] /// - Parameters:

[0974] /// - scale: The scale factor for this view in each dimension.

[0975] /// - anchor: The anchor point about which to scale the view. Defaults to center.

[0976] /// - Returns: A view that scales this view by `scale`.

[0977] @_alwaysEmitIntoClient

[0978] public func scaleEffect(_ scale: Size3D, anchor: UnitPoint3D = .center)

[0979] -> some View

[0980]

[0981] /// Scales this view uniformly by the specified factor.

[0982] ///

[0983] /// The original dimensions of the view are considered to be unchanged by

[0984] /// scaling the contents. To change the dimensions of the view, use a

[0985] /// modifier like `frame()` instead.

[0986] ///

[0987] /// - Parameters:

[0988] /// -s: The scale factor for this view.

[0989] /// -anchor: The anchor point about which to scale the view. Defaults to
center.

[0990] /// - Returns: A view that scales this view by `s` in all dimensions.

[0991] @_disfavoredOverload

[0992] @_alwaysEmitIntoClient

[0993] public func scaleEffect(_ s: CGFloat, anchor: UnitPoint3D = .center)

[0994] -> some View

[0995]

[0996] /// Scales this view by the specified horizontal, vertical, and depth factors.

[0997] ///

[0998] /// The original dimensions of the view are considered to be unchanged by

[0999] /// scaling the contents. To change the dimensions of the view, use a

[1000] /// modifier like `frame()` instead.

[1001] ///

[1002] /// - Parameters:

[1003] /// - x: The horizontal scale factor for this view.

[1004] /// - y: The vertical scale factor for this view.

[1005] /// - z: The depth scale factor for this view.

[1006] /// - anchor: The anchor point about which to scale the view. Defaults to center.

[1007] /// - Returns: A view that scales this view by `x`, `y`, and `z`.

[1008] @_disfavoredOverload

[1009] @_alwaysEmitIntoClient

[1010] public func scaleEffect(
[1011] x: CGFloat = 1.0, y: CGFloat = 1.0, z: CGFloat = 1.0,
[1012] anchor: UnitPoint3D = .center)
[1013] -> some View

[1014]

[1015] /// Applies a 3D transformation to the receiver.

[1016] ///

[1017] /// - Parameters:

```

[1018]    /// - transform: The 3D transformation to apply to the view,
[1019]    ///         interpreting it as a 3D plane in space.
[1020]    /// - Returns: A view that renders transformed according to the provided
[1021]    ///         `transform`
[1022]    @_alwaysEmitIntoClient
[1023]    public func transform3DEffect(_ transform: AffineTransform3D) -> some
View
[1024]    }
[1025]
[1026]    @available(phoneOS 13.0, *)
[1027]    @available(computerOS 10.15, *)
[1028]    @available(mediaOS 13.0, *)
[1029]    @available(wearableOS 6.0, *)
[1030]    @available(mixed reality operating system, introduced: 1.0, deprecated: 1.0,
[1031]    message: "")
[1032]    Use `perspectiveRotationEffect` or a `rotation3DEffect` without a
[1033]    perspective value instead.
[1034]    "")
[1035]    extension View {
[1036]    public func rotation3DEffect(_ angle: Angle,
[1037]    axis: (x: CGFloat, y: CGFloat, z: CGFloat),
[1038]    anchor: UnitPoint = .center, anchorZ: CGFloat = 0,

```

```
[1039]         perspective: CGFloat = 1)
[1040]     -> some View
[1041] }
[1042] ```
[1043] ```swift
[1044] @available(phoneOS, unavailable)
[1045] @available(computerOS, unavailable)
[1046] @available(wearableOS, unavailable)
[1047] @available(mediaOS, unavailable)
[1048] @available(mixed reality operating system 1.0, *)
[1049] extension View {
[1050]     /// Applies effects to this view, while providing access to layout information
[1051]     /// via `GeometryProxy3D`.
[1052]     ///
[1053]     /// You return new effects by calling functions on the first argument
[1054]     /// provided to the `effect` closure. In this example, `ContentView` is
[1055]     /// offset in Z by its own depth, causing its back face to appear where the
[1056]     /// front face of the view was originally located:
[1057]     /// ```swift
[1058]     /// ContentView()
[1059]     ///     .visualEffect3D { content, geometryProxy in
```

```
[1060]    ///    content.offset(z: geometryProxy.size.depth)

[1061]    ///    }

[1062]    ///    ``

[1063]    ///

[1064]    /// - Parameters:

[1065]    /// - effect: A closure that returns the effect to be applied. The first

[1066]    ///    argument provided to the closure is a placeholder representing

[1067]    ///    this view. The second argument is a `GeometryProxy3D`.

[1068]    /// - Returns: A view with the effect applied.

[1069]    public func visualEffect3D(

[1070]        _ effect: @Sendable @escaping (EmptyVisualEffect, GeometryProxy3D) ->
some VisualEffect

[1071]    ) -> some View

[1072]    }

[1073]

[1074]    @available(phoneOS, unavailable)

[1075]    @available(computerOS, unavailable)

[1076]    @available(wearableOS, unavailable)

[1077]    @available(mediaOS, unavailable)

[1078]    @available(mixed reality operating system 1.0, *)

[1079]    extension VisualEffect {

[1080]        /// Brings a view forward in Z by the provided distance in points.
```

[1081] ///

[1082] /// - Parameters:

[1083] /// - distance: The distance to extrude the view forward in Z, in points.

[1084] /// - Returns: An effect that is extruded forward in Z by `distance`.

[1085] public func offset(z: CGFloat) -> some VisualEffect

[1086]

[1087] /// Rotates this view by the specified angle about an axis and anchor point.

[1088] ///

[1089] /// The original dimensions of the view are considered to be unchanged by

[1090] /// rotating the contents.

[1091] ///

[1092] /// For example, to rotate a 3D model by 45° about the Y axis:

[1093] ///

[1094] /// ```

[1095] /// Model3D("robot")

[1096] /// .rotationEffect(.degrees(45), axis: .y)

[1097] /// ```

[1098] ///

[1099] /// - Parameters:

[1100] /// - angle: The angle by which to rotate this view.

[1101] /// - axis: The axis about which to rotate this view.

[1102] /// - anchor: The anchor point about which to rotate this view. Defaults to center.

[1103] /// - Returns: An effect that rotates this view by `angle` about `axis` and `anchor`.

[1104] public func rotation3D(
[1105] _ angle: Angle, axis: RotationAxis3D, anchor: UnitPoint3D = .center)
[1106] -> some VisualEffect
[1107]
[1108] /// Rotates this view by a specific amount.
[1109] ///
[1110] /// The original dimensions of the view are considered to be unchanged by
[1111] /// rotating the contents.
[1112] ///
[1113] /// - Parameters:
[1114] /// - rotation: The rotation by which to rotate the view.
[1115] /// - anchor: The anchor point about which to rotate this view. Defaults to center.
[1116] /// - Returns: An effect that rotates this view by `rotation` about `anchor`.
[1117] public func rotation3D(
[1118] _ rotation: Rotation3D, anchor: UnitPoint3D = .center)
[1119] -> some VisualEffect
[1120]

[1121] /// Rotates this view by the specified angle about an axis and anchor point.

[1122] ///

[1123] /// The original dimensions of the view are considered to be unchanged by

[1124] /// rotating the contents.

[1125] ///

[1126] /// - Parameters:

[1127] /// - angle: The angle by which to rotate this view.

[1128] /// - axis: The axis about which to rotate this view.

[1129] /// - anchor: The anchor point about which to rotate this view. Defaults to center.

[1130] /// - Returns: An effect that rotates the view's rendered output by `angle`

[1131] /// about `axis` and `anchor`.

[1132] public func rotation3D(
[1133] _ angle: Angle, axis: (x: CGFloat, y: CGFloat, z: CGFloat),
[1134] anchor: UnitPoint3D = .center
[1135]) -> some View
[1136]

[1137] /// Rotates the view's rendered output in three dimensions around the given

[1138] /// axis of rotation.

[1139] ///

[1140] /// Use `perspectiveRotation(\_ :axis:anchor:anchorZ:perspective:)` to rotate the

[1141] /// view in three dimensions around the given axis of rotation, and

[1142] /// optionally, position the view at a custom display order and perspective.

[1143] ///

[1144] /// - Parameters:

[1145] /// - angle: The angle at which to rotate the view.

[1146] /// - axis: The `x`, `y` and `z` elements that specify the axis of

[1147] /// rotation.

[1148] /// - anchor: The location with a default of ``UnitPoint/center`` that

[1149] /// defines a point in 3D space about which the rotation is anchored.

[1150] /// - anchorZ: The location with a default of `0` that defines a point in

[1151] /// 3D space about which the rotation is anchored.

[1152] /// - perspective: The relative vanishing point with a default of `1` for

[1153] /// this rotation.

[1154] ///

[1155] /// - Returns: An effect that rotates the view's rendered output in three

[1156] /// dimensions.

[1157] public func perspectiveRotation(_ angle: Angle,

[1158] axis: (x: CGFloat, y: CGFloat, z: CGFloat),

[1159] anchor: UnitPoint3D = .back,

[1160] perspective: CGFloat = 1)

[1161] -> some VisualEffect

[1162]

[1163] /// Scales this view uniformly by the specified size in each dimension.

[1164] ///

[1165] /// The original dimensions of the view are considered to be unchanged by

[1166] /// scaling the contents. To change the dimensions of the view, use a

[1167] /// modifier like `frame()` instead.

[1168] ///

[1169] /// - Parameters:

[1170] /// - scale: The scale factor for this view in each dimension.

[1171] /// - anchor: The anchor point about which to scale the view. Defaults to center.

[1172] /// - Returns: An effect that scales this view by `scale`.

[1173] public func scale(_ scale: Size3D, anchor: UnitPoint3D = .center)

[1174] -> some VisualEffect

[1175]

[1176] /// Scales this view uniformly by the specified factor.

[1177] ///

[1178] /// The original dimensions of the view are considered to be unchanged by

[1179] /// scaling the contents. To change the dimensions of the view, use a

[1180] /// modifier like `frame()` instead.

[1181] ///

[1182] /// - Parameters:

[1183] /// - s: The scale factor for this view.

[1184] /// - anchor: The anchor point about which to scale the view. Defaults to center.

[1185] /// - Returns: An effect that scales this view by `s` in all dimensions.

[1186] @_disfavoredOverload

[1187] public func scale(_ s: CGFloat, anchor: UnitPoint3D = .center)

[1188] -> some VisualEffect

[1189]

[1190] /// Scales this view by the specified horizontal, vertical, and depth factors.

[1191] ///

[1192] /// The original dimensions of the view are considered to be unchanged by

[1193] /// scaling the contents. To change the dimensions of the view, use a

[1194] /// modifier like `frame()` instead.

[1195] ///

[1196] /// - Parameters:

[1197] /// - x: The horizontal scale factor for this view.

[1198] /// - y: The vertical scale factor for this view.

[1199] /// - z: The depth scale factor for this view.

[1200] /// - anchor: The anchor point about which to scale the view. Defaults to center.

[1201] /// - Returns: An effect that scales this view by `x`, `y`, and `z`.

[1202] @_disfavoredOverload

[1203] public func scale(

```
[1204]         x: CGFloat = 1.0, y: CGFloat = 1.0, z: CGFloat = 1.0,
[1205]         anchor: UnitPoint3D = .center)
[1206]     -> some VisualEffect
[1207]
[1208]     /// Applies a 3D transformation to the receiver.
[1209]     ///
[1210]     /// - Parameters:
[1211]     ///   - transform: The 3D transformation to apply to the view,
[1212]     ///     interpreting it as a 3D plane in space.
[1213]     /// - Returns: An effect that renders transformed according to the provided
[1214]     ///   `transform`
[1215]     public func transform3D(_ transform: AffineTransform3D) -> some
VisualEffect
[1216] }
[1217]
[1218] @available(phoneOS 17.0, computerOS 14.0, mediaOS 17.0, wearableOS 10.0, *)
[1219] @available(mixed reality operating system, introduced: 1.0, deprecated: 1.0,
[1220]     message: ""
[1221]     Use `perspectiveRotation` or a `rotation3D` without a perspective value
[1222]     instead.
[1223]     "")
[1224] extension VisualEffect {
```

```
[1225]     public func rotation3D(_ angle: Angle,
[1226]         axis: (x: CGFloat, y: CGFloat, z: CGFloat),
[1227]         anchor: UnitPoint = .center,
[1228]         anchorZ: CGFloat = 0,
[1229]         perspective: CGFloat = 1)
[1230]     -> some VisualEffect
[1231] }
[1232] ```
[1233] ```swift
[1234] /// A normalized 3D point in a view's coordinate space.
[1235] ///
[1236] /// Use a 3D unit point to represent a three-dimensional location in a view
[1237] /// without having to know the view's rendered size. The point stores a value
[1238] /// in each dimension that indicates the fraction of the view's size in that
[1239] /// dimension --- measured from the view's origin --- where the point appears.
[1240] /// For example, you can create a unit point that represents the center of any
[1241] /// view by using the value `0.5` for each dimension:
[1242] ///
[1243] /// let unitPoint = UnitPoint3D(x: 0.5, y: 0.5, z: 0.5)
[1244] ///
[1245] /// > Note: If you need a two-dimensional unit point, use ``UnitPoint`` instead.
```

```
[1246]    ///  
  
[1247]    ///  
[1248]    ///  
[1249]    ///  
[1250]    ///  
  
[1251]    ///  
[1252]    ///  
[1253]    ///  
[1254]    ///  
[1255]    ///  
[1256]    ///  
  
[1257]    ///  
[1258]    ///  
[1259]    ///  
[1260]    ///  
[1261]    ///  
[1262]    ///  
[1263]    ///  
  
[1264]    ///  
[1265]    ///  

```

- [1266] /// like ``zero``, ``center``, or ``topTrailing``. The built-in values
- [1267] /// correspond to common anchor points that you might want to refer to, like
- [1268] /// the center of one of a view's edges.
- [1269] ///
- [1270] /// > Note: A unit point with one or more components outside the range `[0, 1]`
- [1271] /// projects to a point outside of the view.
- [1272] ///
- [1273] /// ### Layout direction
- [1274] ///
- [1275] /// When a person configures their device to use a left-to-right language like
- [1276] /// English, the system places the view's origin in its top-left-back corner,
- [1277] /// with positive x toward the right, positive y toward the bottom of the
- [1278] /// view, and positive z toward the front. In a right-to-left environment, the
- [1279] /// origin moves to the upper-right-back corner, and the positive x direction
- [1280] /// changes to be toward the left. You don't typically need to do anything to
- [1281] /// handle this change, because two-dimensional framework applies the change to
all aspects of the
- [1282] /// system. For example, see the discussion about layout direction in
- [1283] /// ``HorizontalAlignment``.
- [1284] ///
- [1285] /// It's important to test your app for the different locales that you
- [1286] /// distribute your app in. For more information about the localization process,

[1287] /// see <doc://com.apple.documentation/documentation/Xcode/localization>.

[1288] @available(phoneOS, unavailable)

[1289] @available(computerOS, unavailable)

[1290] @available(wearableOS, unavailable)

[1291] @available(mediaOS, unavailable)

[1292] @available(mixed reality operating system 1.0, *)

[1293] @frozen

[1294] public struct UnitPoint3D: Hashable, Codable {

[1295] /// The normalized distance from the origin to the point in the horizontal

[1296] /// direction.

[1297] public var x: CGFloat

[1298]

[1299] /// The normalized distance from the origin to the point in the vertical

[1300] /// dimension.

[1301] public var y: CGFloat

[1302]

[1303] /// The normalized distance from the origin to the point in the depth

[1304] /// dimension.

[1305] public var z: CGFloat

[1306]

[1307] /// Creates a 3D unit point at the origin.

[1308] ///

[1309] /// A view's origin appears in the top-left-back corner in a left-to-right

[1310] /// language environment, with positive x toward the right. It appears in

[1311] /// the top-right-back corner in a right-to-left language, with positive x

[1312] /// toward the left. Positive y is always toward the bottom of the view, and

[1313] /// positive z points toward the front.

[1314] @_alwaysEmitIntoClient

[1315] public init()

[1316]

[1317] /// Creates a 3D unit point with the specified offsets.

[1318] ///

[1319] /// Values outside the range `[0, 1]` project to points outside of a view.

[1320] ///

[1321] /// - Parameters:

[1322] /// - x: The normalized distance from the origin to the point in the

[1323] /// horizontal dimension.

[1324] /// - y: The normalized distance from the origin to the point in the

[1325] /// vertical dimension.

[1326] /// - z: The normalized distance from the origin to the point in the

[1327] /// depth dimension.

[1328] @_alwaysEmitIntoClient

```
[1329]    public init(x: CGFloat, y: CGFloat, z: CGFloat)

[1330]

[1331]    /// A 3D unit point with all components equal to zero.

[1332]    ///

[1333]    /// This point is equivalent to the ``origin``.

[1334]    /// A view's origin appears in the top-left-back corner in a left-to-right

[1335]    /// language environment, with positive x toward the right. It appears in

[1336]    /// the top-right-back corner in a right-to-left language, with positive x

[1337]    /// toward the left. Positive y is always toward the bottom of the view,

[1338]    /// and positive z points toward the front.

[1339]    public static let zero = UnitPoint3D()

[1340]

[1341]    /// A point that's centered in a view.

[1342]    public static let center = UnitPoint3D(x: 0.5, y: 0.5, z: 0.5)

[1343]

[1344]    /// A point that's centered horizontally and in the depth dimension

[1345]    /// on the bottom face of a view.

[1346]    public static let bottom = UnitPoint3D(x: 0.5, y: 1, z: 0.5)

[1347]

[1348]    /// A point that's centered horizontally on the bottom-back edge of a view.

[1349]    public static let bottomBack = UnitPoint3D(x: 0.5, y: 1, z: 0)
```

[1350]

[1351] /// A point that's centered horizontally on the bottom-front edge of a view.

[1352] public static let bottomFront = UnitPoint3D(x: 0.5, y: 1, z: 1)

[1353]

[1354] /// A point that's centered in the depth dimension on the bottom-leading

[1355] /// edge of a view.

[1356] ///

[1357] /// The leading edge appears on the left in a left-to-right language

[1358] /// environment and on the right in a right-to-left environment.

[1359] public static let bottomLeading = UnitPoint3D(x: 0, y: 1, z: 0.5)

[1360]

[1361] /// A point that's in the bottom-leading-back corner of a view.

[1362] ///

[1363] /// The leading edge appears on the left in a left-to-right language

[1364] /// environment and on the right in a right-to-left environment.

[1365] public static let bottomLeadingBack = UnitPoint3D(x: 0, y: 1, z: 0)

[1366]

[1367] /// A point that's in the bottom-leading-front corner of a view.

[1368] ///

[1369] /// The leading edge appears on the left in a left-to-right language

[1370] /// environment and on the right in a right-to-left environment.

[1371] public static let bottomLeadingFront = UnitPoint3D(x: 0, y: 1, z: 1)

[1372]

[1373] /// A point that's centered in the depth dimension on the bottom-trailing

[1374] /// edge of a view.

[1375] ///

[1376] /// The trailing edge appears on the right in a left-to-right language

[1377] /// environment and on the left in a right-to-left environment.

[1378] public static let bottomTrailing = UnitPoint3D(x: 1, y: 1, z: 0.5)

[1379]

[1380] /// A point that's in the bottom-trailing-back corner of a view.

[1381] ///

[1382] /// The trailing edge appears on the right in a left-to-right language

[1383] /// environment and on the left in a right-to-left environment.

[1384] public static let bottomTrailingBack = UnitPoint3D(x: 1, y: 1, z: 0)

[1385]

[1386] /// A point that's in the bottom-trailing-front corner of a view.

[1387] ///

[1388] /// The trailing edge appears on the right in a left-to-right language

[1389] /// environment and on the left in a right-to-left environment.

[1390] public static let bottomTrailingFront = UnitPoint3D(x: 1, y: 1, z: 1)

[1391]

[1392] /// A point that's centered vertically and in the depth dimension

[1393] /// on the leading face of a view.

[1394] ///

[1395] /// The leading edge appears on the left in a left-to-right language

[1396] /// environment and on the right in a right-to-left environment.

[1397] public static let leading = UnitPoint3D(x: 0, y: 0.5, z: 0.5)

[1398]

[1399] /// A point that's centered vertically on the leading-back edge of a view.

[1400] ///

[1401] /// The leading edge appears on the left in a left-to-right language

[1402] /// environment and on the right in a right-to-left environment.

[1403] public static let leadingBack = UnitPoint3D(x: 0, y: 0.5, z: 0)

[1404]

[1405] /// A point that's centered vertically on the leading-front edge of a view.

[1406] ///

[1407] /// The leading edge appears on the left in a left-to-right language

[1408] /// environment and on the right in a right-to-left environment.

[1409] public static let leadingFront = UnitPoint3D(x: 0, y: 0.5, z: 1)

[1410]

[1411] /// A point that's centered horizontally and in the depth dimension

[1412] /// on the top face of a view.

[1413] public static let top = UnitPoint3D(x: 0.5, y: 0, z: 0.5)

[1414]

[1415] /// A point that's centered horizontally on the top-back edge of a view.

[1416] public static let topBack = UnitPoint3D(x: 0.5, y: 0, z: 0)

[1417]

[1418] /// A point that's centered horizontally on the top-front edge of a view.

[1419] public static let topFront = UnitPoint3D(x: 0.5, y: 0, z: 1)

[1420]

[1421] /// A point that's centered in the depth dimension on the top-leading

[1422] /// edge of a view.

[1423] ///

[1424] /// The leading edge appears on the left in a left-to-right language

[1425] /// environment and on the right in a right-to-left environment.

[1426] public static let topLeading = UnitPoint3D(x: 0, y: 0, z: 0.5)

[1427]

[1428] /// A point that's in the top-leading-back corner of a view.

[1429] ///

[1430] /// The leading edge appears on the left in a left-to-right language

[1431] /// environment and on the right in a right-to-left environment.

[1432] public static let topLeadingBack = UnitPoint3D(x: 0, y: 0, z: 0)

[1433]

[1434] /// A point that's in the top-leading-front corner of a view.

[1435] ///

[1436] /// The leading edge appears on the left in a left-to-right language

[1437] /// environment and on the right in a right-to-left environment.

[1438] public static let topLeadingFront = UnitPoint3D(x: 0, y: 0, z: 1)

[1439]

[1440] /// A point that's centered in the depth dimension on the top-trailing

[1441] /// edge of a view.

[1442] ///

[1443] /// The trailing edge appears on the right in a left-to-right language

[1444] /// environment and on the left in a right-to-left environment.

[1445] public static let topTrailing = UnitPoint3D(x: 1, y: 0, z: 0.5)

[1446]

[1447] /// A point that's in the top-trailing-back corner of a view.

[1448] ///

[1449] /// The trailing edge appears on the right in a left-to-right language

[1450] /// environment and on the left in a right-to-left environment.

[1451] public static let topTrailingBack = UnitPoint3D(x: 1, y: 0, z: 0)

[1452]

[1453] /// A point that's in the top-trailing-front corner of a view.

[1454] ///

[1455] /// The trailing edge appears on the right in a left-to-right language

[1456] /// environment and on the left in a right-to-left environment.

[1457] public static let topTrailingFront = UnitPoint3D(x: 1, y: 0, z: 1)

[1458]

[1459] /// A point that's centered vertically and in the depth dimension

[1460] /// on the trailing face of a view.

[1461] ///

[1462] /// The trailing edge appears on the right in a left-to-right language

[1463] /// environment and on the left in a right-to-left environment.

[1464] public static let trailing = UnitPoint3D(x: 1, y: 0.5, z: 0.5)

[1465]

[1466] /// A point that's centered vertically on the trailing-back edge of a view.

[1467] ///

[1468] /// The trailing edge appears on the right in a left-to-right language

[1469] /// environment and on the left in a right-to-left environment.

[1470] public static let trailingBack = UnitPoint3D(x: 1, y: 0.5, z: 0)

[1471]

[1472] /// A point that's centered vertically on the trailing-front edge of a

[1473] /// view.

[1474] ///

[1475] /// The trailing edge appears on the right in a left-to-right language

[1476] /// environment and on the left in a right-to-left environment.

[1477] public static let trailingFront = UnitPoint3D(x: 1, y: 0.5, z: 1)

[1478]

[1479] /// A point that's centered horizontally and vertically

[1480] /// on the front face of a view.

[1481] public static let front = UnitPoint3D(x: 0.5, y: 0.5, z: 1)

[1482]

[1483] /// A point that's centered horizontally and vertically

[1484] /// on the back face of a view.

[1485] public static let back = UnitPoint3D(x: 0.5, y: 0.5, z: 0)

[1486] }

[1487] '''

[1488] As noted earlier, the deprecation of existing rotation3DEffect modifiers on mixed reality operating system, along with an mixed reality operating system -exclusive perspectiveRotationEffect modifier, produces an inconsistency with other platforms. While we believe this is acceptable for now to help limit the scope of this proposal, we do believe it should be remedied in the future by standardizing on modifiers that are available on all platforms where possible. While more thinking is needed here, this is one possible solution: Introduce the perspectiveRotationEffect API and the new rotation3DEffect overloads that do not take a perspective value on all platforms. Note that these modifiers as proposed use the new UnitPoint3D type, which necessitates it being made available on all platforms as well. Deprecate the rotation3DEffect APIs that take a perspective value on all platforms in favor of using perspectiveRotationEffect. On 2D platforms, we would likely use soft deprecation, as using these APIs in those contexts does not meet the bar of being "actively harmful", while we do believe that to be the case on a 3D platform like mixed reality operating system.

[1489] Animating 3D rotations can be tricky to implement correctly. Typically, 3D rotations are best animated using a spherically linear interpolation, or `slerp`, producing a smooth path between two 3D rotations. This requires being able to interpolate between two values in a way that is possible but difficult to implement using two-dimensional framework animation system. The `rotation3DEffect` modifier can implement this animation behavior automatically for most apps; however, for apps that wish to use similar animations themselves with custom 3D transforms, it may be desirable to provide this more directly. It may be possible to have the `Rotation3D` type from the `Spatial` library conform to `Animatable` and implement this behavior, but this has proven difficult to get right in previous attempts. A new `Animatable` type may be required instead. This direction needs more investigation; for now, we are using a custom opaque `Animatable` type for implementing `rotation3DEffect`'s `Animatable` conformance.

[1490] Two-dimensional framework provides a `GeometryEffect` protocol, a kind of `ViewModifier` that can be used to apply any 3x3 `ProjectionTransform` as a function of the view's 2D size. All geometry effects in two-dimensional framework are currently implemented using that protocol. Behind the scenes, two-dimensional framework on mixed reality operating system also has a `GeometryEffect3D` protocol that follows the same design and is how the proposed modifiers like `offset(z:)` and `transform3DEffect` are implemented, with two key differences: The effect returns a 4x3 `AffineTransform3D` instead of a `ProjectionTransform`. The effect is provided a 3D size to allow it to provide transforms that are a function of the view's depth (e.g., a 3D rotation applied about the view's center in all axes). While we considered exposing this protocol publicly as well, we believe this is not necessary with the new `.visualEffect3D` modifier. With this modifier, apps are given even more power than with `GeometryEffect3D`, as they can not only determine the size of a view with the provided `GeometryProxy3D`, but also compute its relative frame or bounds in an ancestor coordinate space like a `ScrollView`. In some examples, the feature not supported by using `VisualEffect`s at this time is the `ignoredByLayout()` modifier on `GeometryEffect` to have a transform effect not affect the view's layout. However, this appears to be quite uncommon to use in most apps, and it could always be added to the `VisualEffect` API later if it proves desirable. While we could also consider exposing `GeometryEffect3D` publicly in the future, it seemed better to encourage use of `.visualEffect3D` for now instead.

(3D LAYOUT) With N301 and mixed reality operating system, we are bringing two-dimensional framework into the 3D world, enabling new kinds of applications and experiences built using 3D content,

gestures, and effects. A common use case for working in a 3D environment is to apply true 3D transforms to content. On mixed reality operating system, this is often used to: Add Z offsets to views to provide a sense of visual hierarchy and emphasis on select elements, even in a 2.5D windowed application, and apply 3D rotations, scales, and offsets to content, such as when manipulating a 3D model with a gesture

[1491] Because two-dimensional framework on mixed reality operating system supports use of 3D content through APIs like ``Model3D`` and ``View`` (proposed separately), as well as 3D scenes like volumetric windows and unbounded Stages, it is also important to be able to apply these 3D transforms relative to a view's 3D size, such as to rotate a 3D model about its center. Behind the scenes, mixed reality operating system already supports applying 3D transforms both to 3D content (via three-dimensional kit) and to 2D views. When a 3D transform is applied to a 2D view, we render that view in a separate mesh and texture through a process known as "separation" baked into Core Animation and three-dimensional kit, which we leverage internally to two-dimensional framework. We aim to expose these capabilities publicly in two-dimensional framework on mixed reality operating systems with a mix of new APIs and new behaviors on existing APIs.

[1492] We propose: A suite of new ``View`` and ``VisualEffect`` modifiers for applying 3D transform effects to a view, a new ``.visualEffect3D`` modifier, analogous to ``.visualEffect`` while providing access to a view's 3D geometry, a mending the existing ``rotation3DEffect`` modifiers to provide a clearer story differentiating perspective rotations from true 3D rotations, and a new ``UnitPoint3D`` type, analogous to ``UnitPoint``, used throughout the above APIs.

[1493] First, we introduce modifiers to perform true 3D transform effects like scales, rotations, and translations on two-dimensional framework views. These are exposed both as ``View`` modifiers and as ``VisualEffect`` modifiers, as with existing transform effect modifiers like ``.scaleEffect()``. The simplest and most commonly used 3D transform effect is to apply a Z offset to a view. These are used throughout mixed reality operating system, such as to present sheets visually in front of its presenting view, or to have controls lift up to meet the user's finger on direct touch hover. For example:

[1494] ````swift Capsule().offset(z: hovering ? 10 : 0)````

[1495] We can also apply true 3D rotations using new variants of the
 `rotation3DEffect()` modifier, which is particularly useful for orienting 3D content:

```
[1496]   ``swift struct Earth: View { @GestureState var rotation: Rotation3D = .identity
var body: some View {
Model3D("earth").rotation3DEffect(rotation).gesture(RotateGesture3D().updating($rotation)
})
```

```
[1497]   }
```

```
[1498] }
```

```
[1499] ``
```

[1500] ****Note****: The `Rotation3D` type is defined in the Spatial framework, as are other
 3D geometry APIs used throughout this proposal like `AffineTransform3D` and `Size3D`.

[1501] The existing `scaleEffect()` modifier also works automatically with 3D content
 when using a uniform scale factor, applying that same scale in the Z axis. This enables
 intuitive behavior in the common case, like with applying a magnify gesture:

```
[1502]   ``swift
```

```
[1503]   struct Earth: View {
```

```
[1504]       @GestureState var scale = 1.0
```

```
[1505]       var body: some View {
```

```
[1506]           Model3D("earth")
```

```
[1507]               .scaleEffect(scale)
```

```
[1508]               .gesture(MagnifyGesture().updating($scale) {
```

```
[1509]                   ...
```

```
[1510]               })
```

```
[1511]       }
```

[1512] }

[1513] ```

[1514] If more granular control is needed, the new `.scaleEffect(x:y:z:anchor:)` modifier can apply non-uniform 3D scale effects as well. Finally, we also add modifiers to apply a custom `AffineTransform3D` 4x3 transform to a view via the `.transform3DEffect()` modifier, analogous to the `.transformEffect()` modifier for applying a 3x3 `ProjectionTransform`. Because these modifiers are also available as `VisualEffect`'s, we can use them within APIs like the `scrollTransition` modifier — for example, to apply 3D effects to views in a `ScrollView`:

[1515] ```swift

[1516] struct ContentView: View {

[1517] var body: some View {

[1518] ScrollView {

[1519] ForEach(City.all) { city in

[1520] CityCard(city)

[1521] .scrollTransition { content, phase in

[1522] content

[1523] .offset(z: phase.isIdentity ? -10 : 0)

[1524] }

[1525] }

[1526] }

[1527] }

[1528] }

[1529] ` ``

[1530] To create custom 3D geometry effects, it is important to be able to read the 3D geometry of the view, such as to apply a transform about a 3D anchor point. The new `.visualEffect` modifier provides access to a `GeometryProxy` within its closure. In some examples, this includes the view's 2D size. To address this, introduce a `.visualEffect3D` modifier, which behaves like `.visualEffect` but instead provides a `GeometryProxy3D` (as proposed[here](https://github.pie.apple.com/TDG/two-dimensional-framework-Evolution/pull/14)). This enables more advanced effects, like a recessing effect that is proportional to the depth of a `Model3D`:

[1531] ` `` swift

[1532] struct RecedingModelExample: View {

[1533] @State var receded = false

[1534] var body: some View {

[1535] Model3D("starship")

[1536] .visualEffect3D { geometry, content in

[1537] content

[1538] .offset(z: receded ? -geometry.depth : 0)

[1539] .scaleEffect(recessed ? 0.5 : 1.0)

[1540] }

[1541] Toggle("Receded", isOn: \$receded)

[1542] }

[1543] }

[1544] ` ``

[1545] As noted earlier, we propose using the `rotation3DEffect` name to generally represent a true 3D rotation on mixed reality operating system, as this most accurately reflects the behavior expressed by the modifier name. To model the behavior of a true 3D rotation, we introduce new overloads with 3D types like `Rotation3D`:

[1546] ```swift

[1547] extension View {

[1548] public func rotation3DEffect(

[1549] _ angle: Angle, axis: RotationAxis3D, anchor: UnitPoint3D = .center)

[1550] -> some View

[1551] public func rotation3DEffect(

[1552] _ rotation: Rotation3D, anchor: UnitPoint3D = .center)

[1553] -> some View

[1554] public func rotation3DEffect(

[1555] _ angle: Angle, axis: (x: CGFloat, y: CGFloat, z: CGFloat),

[1556] anchor: UnitPoint3D = .center

[1557]) -> some View

[1558] }

[1559] ```

[1560] Note that unlike the existing `rotation3DEffect` modifiers, these overloads do **not** take a `perspective` value, as the concept does not have a sensible analog with a true 3D rotation. However, there are valid use cases on mixed reality operating system for preserving the existing behavior of `rotation3DEffect` on other platforms, where a **perspective** rotation is applied to the view as a 2D projection, without creating true 3D geometry. To support these use cases, we add a new `perspectiveRotationEffect` modifier:

[1561] ```swift

[1562] extension View {

[1563] public func perspectiveRotationEffect(_ angle: Angle,

[1564] axis: (x: CGFloat, y: CGFloat, z: CGFloat),

[1565] anchor: UnitPoint3D = .back, perspective: CGFloat = 1)

[1566] -> some View

[1567] }

[1568] ```

[1569] To help further disambiguate between perspective rotations and true 3D rotations, we propose hard-deprecating the existing `rotation3DEffect` modifiers on mixed reality operating system and directing developers to use `perspectiveRotationEffect` instead. In some examples, this affects apps that specify a `perspective` value or `anchor` explicitly, however: when only an angle and axis are specified, the new `rotation3DEffect` overload is preferred instead when compiling for the mixed reality operating system SDK and automatically results in applying a true 3D rotation:

[1570] ```swift

[1571] SomeView()

[1572] .rotation3DEffect(.degrees(45), axis: (x: 1, y: 0, z: 0))

[1573] ```

[1574] Note that this approach does not propose any changes to the 3D rotation APIs for other platforms at this time. See the *Future directions* section for additional thoughts on that subject.`

[1575] To support the above APIs, we introduce a new `UnitPoint3D` type, serving as a 3D analog to the existing `UnitPoint` type. `UnitPoint3D` provides static members for common points, with additional variants to express the front, back, and center of a given 2D

point. For example, ``UnitPoint3D.topLeadingBack`` represents the point at the back of the view in Z and the top-leading corner in X and Y, while ``UnitPoint3D.topLeading`` instead represents the center of the view in Z. Because there is some naming overlap with ``UnitPoint``'s static members, some care must be taken with modifiers that can in some places take a ``UnitPoint3D`` and in others a ``UnitPoint``. For instance, the new ``scaleEffect(x:y:z:anchor:)`` modifier takes a ``UnitPoint3D``, and is marked as ``@_disfavoredOverload`` to ensure that writing code like this:

[1576] ````swift`

[1577] `Color.red.scaleEffect(x: 2, y: 1, anchor: .topLeading)`

[1578] `````

[1579] Biases towards using the existing modifier that takes a ``UnitPoint`` and no ``z`` parameter, preserving existing semantics and expectations for apps that are also built for other platforms.

[1580] ````swift`

[1581] `@available(phoneOS, unavailable)`

[1582] `@available(computerOS, unavailable)`

[1583] `@available(wearableOS, unavailable)`

[1584] `@available(mediaOS, unavailable)`

[1585] `@available(mixed reality operating system 1.0, *)`

[1586] `extension View {`

[1587] `/// Brings a view forward in Z by the provided distance in points.`

[1588] `///`

[1589] `/// - Parameters:`

[1590] `///     - distance: The distance to extrude the view forward in Z, in points.`

[1591] /// - Returns: A view that is extruded forward in Z by `distance`.

[1592] @_alwaysEmitIntoClient

[1593] public func offset(z: CGFloat) -> some View

[1594] /// Rotates this view by the specified angle about an axis and anchor point.

[1595] ///

[1596] /// The original dimensions of the view are considered to be unchanged by

[1597] /// rotating the contents.

[1598] ///

[1599] /// For example, to rotate a 3D model by 45° about the Y axis:

[1600] ///

[1601] /// ```

[1602] /// Model3D("robot")

[1603] /// .rotationEffect(.degrees(45), axis: .y)

[1604] /// ```

[1605] ///

[1606] /// - Parameters:

[1607] /// - angle: The angle by which to rotate this view.

[1608] /// - axis: The axis about which to rotate this view.

[1609] /// - anchor: The anchor point about which to rotate this view. Defaults to
center.

[1610] /// - Returns: A view that rotates this view by `angle` about `axis` and `anchor`.

[1611] @_alwaysEmitIntoClient

```
[1612]     public func rotation3DEffect(  
[1613]         _ angle: Angle, axis: RotationAxis3D, anchor: UnitPoint3D = .center)  
[1614]     -> some View  
[1615]     /// Rotates this view by a specific amount.  
[1616]     ///  
[1617]     /// The original dimensions of the view are considered to be unchanged by  
[1618]     /// rotating the contents.  
[1619]     ///  
[1620]     /// For example, to rotate a 3D model by 45° about the Y axis:  
[1621]     ///  
[1622]     /// ```  
[1623]     /// Model3D("robot")  
[1624]     ///     .rotationEffect(Rotation3D(angle: .degrees(45), axis: .y))  
[1625]     /// ```  
[1626]     ///  
[1627]     /// - Parameters:  
[1628]     ///     - rotation: The rotation by which to rotate the view.  
[1629]     ///     - anchor: The anchor point about which to rotate this view. Defaults to  
[1629]     center.  
[1630]     /// - Returns: A view that rotates this view by `rotation` about `anchor`.  
[1631]     @_alwaysEmitIntoClient  
[1632]     public func rotation3DEffect(  

```

```
[1633]         _rotation: Rotation3D, anchor: UnitPoint3D = .center)

[1634]     -> some View

[1635]     /// Rotates this view by the specified angle about an axis and anchor point.

[1636]     ///

[1637]     /// The original dimensions of the view are considered to be unchanged by

[1638]     /// rotating the contents.

[1639]     ///

[1640]     /// For example, to rotate a 3D model by 45° about the Y axis:

[1641]     ///

[1642]     /// ```

[1643]     /// Model3D("robot")

[1644]     ///     .rotation3DEffect(.degrees(45), axis: .y)

[1645]     /// ```

[1646]     ///

[1647]     /// - Parameters:

[1648]     ///   - angle: The angle by which to rotate this view.

[1649]     ///   - axis: The axis about which to rotate this view.

[1650]     ///   - anchor: The anchor point about which to rotate this view. Defaults to
center.

[1651]     /// - Returns: A view that rotates this view by `angle` about `axis` and `anchor`.

[1652]     @_alwaysEmitIntoClient

[1653]     public func rotation3DEffect(
```

```
[1654]     _angle: Angle, axis: (x: CGFloat, y: CGFloat, z: CGFloat),
[1655]     anchor: UnitPoint3D = .center
[1656] ) -> some View
[1657]     /// Rotates this view's rendered output in three dimensions around the given
[1658]     /// axis of rotation.
[1659]     ///
[1660]     /// Use `perspectiveRotationEffect(_:axis:anchor:anchorZ:perspective:)` to
[1661]     /// rotate the view in three dimensions around the given axis of rotation,
[1662]     /// and optionally, position the view at a custom display order and
[1663]     /// perspective.
[1664]     ///
[1665]     /// In the example below, the text is rotated 45° about the `y` axis,
[1666]     /// front-most (the default `zIndex`) and default `perspective` (`1`):
[1667]     ///
[1668]     /// Text("Rotation by passing an angle in degrees")
[1669]     ///     .perspectiveRotationEffect(.degrees(45), axis: (x: 0, y: 1, z: 0))
[1670]     ///     .border(Color.gray)
[1671]     ///
[1672]     /// ![A screenshot showing the rotation of text 45 degrees about the
[1673]     /// y-axis.](two-dimensional framework-View-rotation3DEffect.png)
[1674]     ///
```

[1675] /// - Parameters:

[1676] /// - angle: The angle at which to rotate the view.

[1677] /// - axis: The `x`, `y` and `z` elements that specify the axis of

[1678] /// rotation.

[1679] /// - anchor: The location with a default of ``UnitPoint/center`` that

[1680] /// defines a point in 3D space about which the rotation is anchored.

[1681] /// - anchorZ: The location with a default of `0` that defines a point in

[1682] /// 3D space about which the rotation is anchored.

[1683] /// - perspective: The relative vanishing point with a default of `1` for

[1684] /// this rotation.

[1685] @_alwaysEmitIntoClient

[1686] public func perspectiveRotationEffect(_ angle: Angle,

[1687] axis: (x: CGFloat, y: CGFloat, z: CGFloat),

[1688] anchor: UnitPoint3D = .back, perspective: CGFloat = 1)

[1689] -> some View

[1690] /// Scales this view uniformly by the specified size in each dimension.

[1691] ///

[1692] /// The original dimensions of the view are considered to be unchanged by

[1693] /// scaling the contents. To change the dimensions of the view, use a

[1694] /// modifier like `frame()` instead.

[1695] ///

[1696] /// - Parameters:

[1697] /// - scale: The scale factor for this view in each dimension.

[1698] /// - anchor: The anchor point about which to scale the view. Defaults to center.

[1699] /// - Returns: A view that scales this view by `scale`.

[1700] @_alwaysEmitIntoClient

[1701] public func scaleEffect(_ scale: Size3D, anchor: UnitPoint3D = .center)

[1702] -> some View

[1703] /// Scales this view uniformly by the specified factor.

[1704] ///

[1705] /// The original dimensions of the view are considered to be unchanged by

[1706] /// scaling the contents. To change the dimensions of the view, use a

[1707] /// modifier like `frame()` instead.

[1708] ///

[1709] /// - Parameters:

[1710] /// - s: The scale factor for this view.

[1711] /// - anchor: The anchor point about which to scale the view. Defaults to center.

[1712] /// - Returns: A view that scales this view by `s` in all dimensions.

[1713] @_disfavoredOverload

[1714] @_alwaysEmitIntoClient

[1715] public func scaleEffect(_ s: CGFloat, anchor: UnitPoint3D = .center)

[1716] -> some View

[1717] /// Scales this view by the specified horizontal, vertical, and depth factors.

[1718] ///

[1719] /// The original dimensions of the view are considered to be unchanged by

[1720] /// scaling the contents. To change the dimensions of the view, use a

[1721] /// modifier like `frame()` instead.

[1722] ///

[1723] /// - Parameters:

[1724] /// -x: The horizontal scale factor for this view.

[1725] /// -y: The vertical scale factor for this view.

[1726] /// -z: The depth scale factor for this view.

[1727] /// -anchor: The anchor point about which to scale the view. Defaults to center.

[1728] /// - Returns: A view that scales this view by `x`, `y`, and `z`.

[1729] @_disfavoredOverload

[1730] @_alwaysEmitIntoClient

[1731] public func scaleEffect(
[1732] x: CGFloat = 1.0, y: CGFloat = 1.0, z: CGFloat = 1.0,
[1733] anchor: UnitPoint3D = .center)
[1734] -> some View

[1735] /// Applies a 3D transformation to the receiver.

[1736] ///

[1737] /// - Parameters:

[1738] /// - transform: The 3D transformation to apply to the view,

[1739] /// interpreting it as a 3D plane in space.

[1740] /// - Returns: A view that renders transformed according to the provided

[1741] /// `transform`

[1742] @_alwaysEmitIntoClient

[1743] public func transform3DEffect(_ transform: AffineTransform3D) -> some
View

[1744] }

[1745] @available(phoneOS 13.0, *)

[1746] @available(computerOS 10.15, *)

[1747] @available(mediaOS 13.0, *)

[1748] @available(wearableOS 6.0, *)

[1749] @available(mixed reality operating system, introduced: 1.0, deprecated: 1.0,

[1750] message: ""

[1751] Use `perspectiveRotationEffect` or a `rotation3DEffect` without a

[1752] perspective value instead.

[1753] """)

[1754] extension View {

[1755] public func rotation3DEffect(_ angle: Angle,

[1756] axis: (x: CGFloat, y: CGFloat, z: CGFloat),

[1757] anchor: UnitPoint = .center, anchorZ: CGFloat = 0,

```
[1758]         perspective: CGFloat = 1)
[1759]     -> some View
[1760] }
[1761] ```
[1762] ```swift
[1763] @available(phoneOS, unavailable)
[1764] @available(computerOS, unavailable)
[1765] @available(wearableOS, unavailable)
[1766] @available(mediaOS, unavailable)
[1767] @available(mixed reality operating system 1.0, *)
[1768] extension View {
[1769]     /// Applies effects to this view, while providing access to layout information
[1770]     /// via `GeometryProxy3D`.
[1771]     ///
[1772]     /// You return new effects by calling functions on the first argument
[1773]     /// provided to the `effect` closure. In this example, `ContentView` is
[1774]     /// offset in Z by its own depth, causing its back face to appear where the
[1775]     /// front face of the view was originally located:
[1776]     /// ```swift
[1777]     /// ContentView()
[1778]     ///     .visualEffect3D { content, geometryProxy in
```

```
[1779]    ///    content.offset(z: geometryProxy.size.depth)

[1780]    ///    }

[1781]    ///    ``

[1782]    ///

[1783]    /// - Parameters:

[1784]    /// - effect: A closure that returns the effect to be applied. The first

[1785]    ///    argument provided to the closure is a placeholder representing

[1786]    ///    this view. The second argument is a `GeometryProxy3D`.

[1787]    /// - Returns: A view with the effect applied.

[1788]    public func visualEffect3D(

[1789]        _ effect: @Sendable @escaping (EmptyVisualEffect, GeometryProxy3D) ->

some VisualEffect

[1790]    ) -> some View

[1791]    }

[1792]    @available(phoneOS, unavailable)

[1793]    @available(computerOS, unavailable)

[1794]    @available(wearableOS, unavailable)

[1795]    @available(mediaOS, unavailable)

[1796]    @available(mixed reality operating system 1.0, *)

[1797]    extension VisualEffect {

[1798]        /// Brings a view forward in Z by the provided distance in points.

[1799]        ///
```

[1800] /// - Parameters:

[1801] /// - distance: The distance to extrude the view forward in Z, in points.

[1802] /// - Returns: An effect that is extruded forward in Z by `distance`.

[1803] public func offset(z: CGFloat) -> some VisualEffect

[1804]

[1805] /// Rotates this view by the specified angle about an axis and anchor point.

[1806] ///

[1807] /// The original dimensions of the view are considered to be unchanged by

[1808] /// rotating the contents.

[1809] ///

[1810] /// For example, to rotate a 3D model by 45° about the Y axis:

[1811] ///

[1812] /// ```

[1813] /// Model3D("robot")

[1814] /// .rotationEffect(.degrees(45), axis: .y)

[1815] /// ```

[1816] ///

[1817] /// - Parameters:

[1818] /// - angle: The angle by which to rotate this view.

[1819] /// - axis: The axis about which to rotate this view.

[1820] /// - anchor: The anchor point about which to rotate this view. Defaults to center.

[1821] /// - Returns: An effect that rotates this view by `angle` about `axis` and
`anchor`.

[1822] public func rotation3D(
[1823] _angle: Angle, axis: RotationAxis3D, anchor: UnitPoint3D = .center)
[1824] -> some VisualEffect

[1825] /// Rotates this view by a specific amount.

[1826] ///

[1827] /// The original dimensions of the view are considered to be unchanged by
[1828] /// rotating the contents.

[1829] ///

[1830] /// - Parameters:

[1831] /// - rotation: The rotation by which to rotate the view.

[1832] /// - anchor: The anchor point about which to rotate this view. Defaults to
center.

[1833] /// - Returns: An effect that rotates this view by `rotation` about `anchor`.

[1834] public func rotation3D(
[1835] _rotation: Rotation3D, anchor: UnitPoint3D = .center)
[1836] -> some VisualEffect

[1837] /// Rotates this view by the specified angle about an axis and anchor point.

[1838] ///

[1839] /// The original dimensions of the view are considered to be unchanged by
[1840] /// rotating the contents.

```

[1841]    ///
[1842]    /// - Parameters:
[1843]    /// - angle: The angle by which to rotate this view.
[1844]    /// - axis: The axis about which to rotate this view.
[1845]    /// - anchor: The anchor point about which to rotate this view. Defaults to
center.
[1846]    /// - Returns: An effect that rotates the view's rendered output by `angle`
[1847]    /// about `axis` and `anchor`.
[1848]    public func rotation3D(
[1849]        _ angle: Angle, axis: (x: CGFloat, y: CGFloat, z: CGFloat),
[1850]        anchor: UnitPoint3D = .center
[1851]    ) -> some View
[1852]    /// Rotates the view's rendered output in three dimensions around the given
[1853]    /// axis of rotation.
[1854]    ///
[1855]    /// Use `perspectiveRotation(_:_axis:_anchor:_anchorZ:_perspective:_)` to rotate the
[1856]    /// view in three dimensions around the given axis of rotation, and
[1857]    /// optionally, position the view at a custom display order and perspective.
[1858]    ///
[1859]    /// - Parameters:
[1860]    /// - angle: The angle at which to rotate the view.
[1861]    /// - axis: The `x`, `y` and `z` elements that specify the axis of

```

[1862] /// rotation.

[1863] /// - anchor: The location with a default of ``UnitPoint/center`` that

[1864] /// defines a point in 3D space about which the rotation is anchored.

[1865] /// - anchorZ: The location with a default of `0` that defines a point in

[1866] /// 3D space about which the rotation is anchored.

[1867] /// - perspective: The relative vanishing point with a default of `1` for

[1868] /// this rotation.

[1869] ///

[1870] /// - Returns: An effect that rotates the view's rendered output in three

[1871] /// dimensions.

[1872] public func perspectiveRotation(_ angle: Angle,

[1873] axis: (x: CGFloat, y: CGFloat, z: CGFloat),

[1874] anchor: UnitPoint3D = .back,

[1875] perspective: CGFloat = 1)

[1876] -> some VisualEffect

[1877] /// Scales this view uniformly by the specified size in each dimension.

[1878] ///

[1879] /// The original dimensions of the view are considered to be unchanged by

[1880] /// scaling the contents. To change the dimensions of the view, use a

[1881] /// modifier like `frame()` instead.

[1882] ///

[1883] /// - Parameters:

[1884] /// - scale: The scale factor for this view in each dimension.

[1885] /// - anchor: The anchor point about which to scale the view. Defaults to center.

[1886] /// - Returns: An effect that scales this view by `scale`.

[1887] public func scale(_ scale: Size3D, anchor: UnitPoint3D = .center)

[1888] -> some VisualEffect

[1889] /// Scales this view uniformly by the specified factor.

[1890] ///

[1891] /// The original dimensions of the view are considered to be unchanged by

[1892] /// scaling the contents. To change the dimensions of the view, use a

[1893] /// modifier like `frame()` instead.

[1894] ///

[1895] /// - Parameters:

[1896] /// - s: The scale factor for this view.

[1897] /// - anchor: The anchor point about which to scale the view. Defaults to center.

[1898] /// - Returns: An effect that scales this view by `s` in all dimensions.

[1899] @_disfavoredOverload

[1900] public func scale(_ s: CGFloat, anchor: UnitPoint3D = .center)

[1901] -> some VisualEffect

[1902] /// Scales this view by the specified horizontal, vertical, and depth factors.

[1903] ///

[1904] /// The original dimensions of the view are considered to be unchanged by

[1905] /// scaling the contents. To change the dimensions of the view, use a

[1906] /// modifier like `frame()` instead.

[1907] ///

[1908] /// - Parameters:

[1909] /// - x: The horizontal scale factor for this view.

[1910] /// - y: The vertical scale factor for this view.

[1911] /// - z: The depth scale factor for this view.

[1912] /// - anchor: The anchor point about which to scale the view. Defaults to center.

[1913] /// - Returns: An effect that scales this view by `x`, `y`, and `z`.

[1914] @_disfavoredOverload

[1915] public func scale(
[1916] x: CGFloat = 1.0, y: CGFloat = 1.0, z: CGFloat = 1.0,
[1917] anchor: UnitPoint3D = .center)
[1918] -> some VisualEffect

[1919] /// Applies a 3D transformation to the receiver.

[1920] ///

[1921] /// - Parameters:

[1922] /// - transform: The 3D transformation to apply to the view,
[1923] /// interpreting it as a 3D plane in space.

```

[1924]    /// - Returns: An effect that renders transformed according to the provided
[1925]    ///    `transform`
[1926]    public func transform3D(_ transform: AffineTransform3D) -> some
VisualEffect
[1927]    }
[1928]    @available(phoneOS 17.0, computerOS 14.0, mediaOS 17.0, wearableOS 10.0, *)
[1929]    @available(mixed reality operating system, introduced: 1.0, deprecated: 1.0,
[1930]    message: "")
[1931]    Use `perspectiveRotation` or a `rotation3D` without a perspective value
[1932]    instead.
[1933]    "")
[1934]    extension VisualEffect {
[1935]    public func rotation3D(_ angle: Angle,
[1936]    axis: (x: CGFloat, y: CGFloat, z: CGFloat),
[1937]    anchor: UnitPoint = .center,
[1938]    anchorZ: CGFloat = 0,
[1939]    perspective: CGFloat = 1)
[1940]    -> some VisualEffect
[1941]    }
[1942]    ``
[1943]
[1944]    ``swift

```

[1945] /// A normalized 3D point in a view's coordinate space.

[1946] ///

[1947] /// Use a 3D unit point to represent a three-dimensional location in a view

[1948] /// without having to know the view's rendered size. The point stores a value

[1949] /// in each dimension that indicates the fraction of the view's size in that

[1950] /// dimension --- measured from the view's origin --- where the point appears.

[1951] /// For example, you can create a unit point that represents the center of any

[1952] /// view by using the value `0.5` for each dimension:

[1953] ///

[1954] /// let unitPoint = UnitPoint3D(x: 0.5, y: 0.5, z: 0.5)

[1955] ///

[1956] /// > Note: If you need a two-dimensional unit point, use ``UnitPoint`` instead.

[1957] ///

[1958] /// To project the unit point into the rendered view's coordinate space,

[1959] /// multiply each component of the unit point with the corresponding

[1960] /// component of the view's size:

[1961] ///

[1962] /// let projectedPoint = Point3D(
[1963] /// x: unitPoint.x * size.width,
[1964] /// y: unitPoint.y * size.height,
[1965] /// z: unitPoint.z * size.depth

[1966] ///
[1967] ///
[1968] /// You can perform this calculation yourself if you happen to know a view's
[1969] /// size, or if you want to use a unit point for some custom purpose, but
[1970] /// two-dimensional framework typically does this for you to carry out operations
that
[1971] /// you request, like when you rotate a view with the
[1972] /// ``View/rotation3DEffect(\_:anchor:)`` modifier and indicate the anchor point
[1973] /// that you want to rotate the view around.
[1974] ///
[1975] /// You can create custom unit points with explicit values, like the example
[1976] /// above, or you can use one of the built-in unit points that two-dimensional
framework provides,
[1977] /// like ``zero``, ``center``, or ``topTrailing``. The built-in values
[1978] /// correspond to common anchor points that you might want to refer to, like
[1979] /// the center of one of a view's edges.
[1980] ///
[1981] /// > Note: A unit point with one or more components outside the range `[0, 1]`
[1982] /// projects to a point outside of the view.
[1983] ///
[1984] /// ### Layout direction
[1985] ///

[1986] /// When a person configures their device to use a left-to-right language like

[1987] /// English, the system places the view's origin in its top-left-back corner,

[1988] /// with positive x toward the right, positive y toward the bottom of the

[1989] /// view, and positive z toward the front. In a right-to-left environment, the

[1990] /// origin moves to the upper-right-back corner, and the positive x direction

[1991] /// changes to be toward the left. You don't typically need to do anything to

[1992] /// handle this change, because two-dimensional framework applies the change to

all aspects of the

[1993] /// system. For example, see the discussion about layout direction in

[1994] /// ``HorizontalAlignment``.

[1995] ///

[1996] /// It's important to test your app for the different locales that you

[1997] /// distribute your app in. For more information about the localization process,

[1998] /// see <doc://com.apple.documentation/documentation/Xcode/localization>.

[1999] @available(phoneOS, unavailable)

[2000] @available(computerOS, unavailable)

[2001] @available(wearableOS, unavailable)

[2002] @available(mediaOS, unavailable)

[2003] @available(mixed reality operating system 1.0, *)

[2004] @frozen

[2005] public struct UnitPoint3D: Hashable, Codable {

[2006] /// The normalized distance from the origin to the point in the horizontal

[2007] /// direction.

[2008] public var x: CGFloat

[2009]

[2010] /// The normalized distance from the origin to the point in the vertical

[2011] /// dimension.

[2012] public var y: CGFloat

[2013]

[2014] /// The normalized distance from the origin to the point in the depth

[2015] /// dimension.

[2016] public var z: CGFloat

[2017]

[2018] /// Creates a 3D unit point at the origin.

[2019] ///

[2020] /// A view's origin appears in the top-left-back corner in a left-to-right

[2021] /// language environment, with positive x toward the right. It appears in

[2022] /// the top-right-back corner in a right-to-left language, with positive x

[2023] /// toward the left. Positive y is always toward the bottom of the view, and

[2024] /// positive z points toward the front.

[2025] @_alwaysEmitIntoClient

[2026] public init()

[2027] /// Creates a 3D unit point with the specified offsets.

[2028] ///

[2029] /// Values outside the range `[0, 1]` project to points outside of a view.

[2030] ///

[2031] /// - Parameters:

[2032] /// - x: The normalized distance from the origin to the point in the

[2033] /// horizontal dimension.

[2034] /// - y: The normalized distance from the origin to the point in the

[2035] /// vertical dimension.

[2036] /// - z: The normalized distance from the origin to the point in the

[2037] /// depth dimension.

[2038] @_alwaysEmitIntoClient

[2039] public init(x: CGFloat, y: CGFloat, z: CGFloat)

[2040]

[2041] /// A 3D unit point with all components equal to zero.

[2042] ///

[2043] /// This point is equivalent to the ``origin``.

[2044] /// A view's origin appears in the top-left-back corner in a left-to-right

[2045] /// language environment, with positive x toward the right. It appears in

[2046] /// the top-right-back corner in a right-to-left language, with positive x

[2047] /// toward the left. Positive y is always toward the bottom of the view,

[2048] /// and positive z points toward the front.

[2049] public static let zero = UnitPoint3D()
[2050] /// A point that's centered in a view.
[2051] public static let center = UnitPoint3D(x: 0.5, y: 0.5, z: 0.5)
[2052] /// A point that's centered horizontally and in the depth dimension
[2053] /// on the bottom face of a view.
[2054] public static let bottom = UnitPoint3D(x: 0.5, y: 1, z: 0.5)
[2055] /// A point that's centered horizontally on the bottom-back edge of a view.
[2056] public static let bottomBack = UnitPoint3D(x: 0.5, y: 1, z: 0)
[2057] /// A point that's centered horizontally on the bottom-front edge of a view.
[2058] public static let bottomFront = UnitPoint3D(x: 0.5, y: 1, z: 1)
[2059] /// A point that's centered in the depth dimension on the bottom-leading
[2060] /// edge of a view.
[2061] ///
[2062] /// The leading edge appears on the left in a left-to-right language
[2063] /// environment and on the right in a right-to-left environment.
[2064] public static let bottomLeading = UnitPoint3D(x: 0, y: 1, z: 0.5)
[2065] /// A point that's in the bottom-leading-back corner of a view.
[2066] ///
[2067] /// The leading edge appears on the left in a left-to-right language
[2068] /// environment and on the right in a right-to-left environment.
[2069] public static let bottomLeadingBack = UnitPoint3D(x: 0, y: 1, z: 0)

[2070] /// A point that's in the bottom-leading-front corner of a view.

[2071] ///

[2072] /// The leading edge appears on the left in a left-to-right language

[2073] /// environment and on the right in a right-to-left environment.

[2074] public static let bottomLeadingFront = UnitPoint3D(x: 0, y: 1, z: 1)

[2075] /// A point that's centered in the depth dimension on the bottom-trailing

[2076] /// edge of a view.

[2077] ///

[2078] /// The trailing edge appears on the right in a left-to-right language

[2079] /// environment and on the left in a right-to-left environment.

[2080] public static let bottomTrailing = UnitPoint3D(x: 1, y: 1, z: 0.5)

[2081] /// A point that's in the bottom-trailing-back corner of a view.

[2082] ///

[2083] /// The trailing edge appears on the right in a left-to-right language

[2084] /// environment and on the left in a right-to-left environment.

[2085] public static let bottomTrailingBack = UnitPoint3D(x: 1, y: 1, z: 0)

[2086] /// A point that's in the bottom-trailing-front corner of a view.

[2087] ///

[2088] /// The trailing edge appears on the right in a left-to-right language

[2089] /// environment and on the left in a right-to-left environment.

[2090] public static let bottomTrailingFront = UnitPoint3D(x: 1, y: 1, z: 1)

- [2091] /// A point that's centered vertically and in the depth dimension
- [2092] /// on the leading face of a view.
- [2093] ///
- [2094] /// The leading edge appears on the left in a left-to-right language
- [2095] /// environment and on the right in a right-to-left environment.
- [2096] public static let leading = UnitPoint3D(x: 0, y: 0.5, z: 0.5)
- [2097] /// A point that's centered vertically on the leading-back edge of a view.
- [2098] ///
- [2099] /// The leading edge appears on the left in a left-to-right language
- [2100] /// environment and on the right in a right-to-left environment.
- [2101] public static let leadingBack = UnitPoint3D(x: 0, y: 0.5, z: 0)
- [2102] /// A point that's centered vertically on the leading-front edge of a view.
- [2103] ///
- [2104] /// The leading edge appears on the left in a left-to-right language
- [2105] /// environment and on the right in a right-to-left environment.
- [2106] public static let leadingFront = UnitPoint3D(x: 0, y: 0.5, z: 1)
- [2107] /// A point that's centered horizontally and in the depth dimension
- [2108] /// on the top face of a view.
- [2109] public static let top = UnitPoint3D(x: 0.5, y: 0, z: 0.5)
- [2110] /// A point that's centered horizontally on the top-back edge of a view.
- [2111] public static let topBack = UnitPoint3D(x: 0.5, y: 0, z: 0)

- [2112] /// A point that's centered horizontally on the top-front edge of a view.
- [2113] public static let topFront = UnitPoint3D(x: 0.5, y: 0, z: 1)
- [2114] /// A point that's centered in the depth dimension on the top-leading
- [2115] /// edge of a view.
- [2116] ///
- [2117] /// The leading edge appears on the left in a left-to-right language
- [2118] /// environment and on the right in a right-to-left environment.
- [2119] public static let topLeading = UnitPoint3D(x: 0, y: 0, z: 0.5)
- [2120] /// A point that's in the top-leading-back corner of a view.
- [2121] ///
- [2122] /// The leading edge appears on the left in a left-to-right language
- [2123] /// environment and on the right in a right-to-left environment.
- [2124] public static let topLeadingBack = UnitPoint3D(x: 0, y: 0, z: 0)
- [2125] /// A point that's in the top-leading-front corner of a view.
- [2126] ///
- [2127] /// The leading edge appears on the left in a left-to-right language
- [2128] /// environment and on the right in a right-to-left environment.
- [2129] public static let topLeadingFront = UnitPoint3D(x: 0, y: 0, z: 1)
- [2130] /// A point that's centered in the depth dimension on the top-trailing
- [2131] /// edge of a view.
- [2132] ///

- [2133] /// The trailing edge appears on the right in a left-to-right language
- [2134] /// environment and on the left in a right-to-left environment.
- [2135] public static let topTrailing = UnitPoint3D(x: 1, y: 0, z: 0.5)
- [2136] /// A point that's in the top-trailing-back corner of a view.
- [2137] ///
- [2138] /// The trailing edge appears on the right in a left-to-right language
- [2139] /// environment and on the left in a right-to-left environment.
- [2140] public static let topTrailingBack = UnitPoint3D(x: 1, y: 0, z: 0)
- [2141] /// A point that's in the top-trailing-front corner of a view.
- [2142] ///
- [2143] /// The trailing edge appears on the right in a left-to-right language
- [2144] /// environment and on the left in a right-to-left environment.
- [2145] public static let topTrailingFront = UnitPoint3D(x: 1, y: 0, z: 1)
- [2146] /// A point that's centered vertically and in the depth dimension
- [2147] /// on the trailing face of a view.
- [2148] ///
- [2149] /// The trailing edge appears on the right in a left-to-right language
- [2150] /// environment and on the left in a right-to-left environment.
- [2151] public static let trailing = UnitPoint3D(x: 1, y: 0.5, z: 0.5)
- [2152]
- [2153] /// A point that's centered vertically on the trailing-back edge of a view.

[2154] ///

[2155] /// The trailing edge appears on the right in a left-to-right language

[2156] /// environment and on the left in a right-to-left environment.

[2157] public static let trailingBack = UnitPoint3D(x: 1, y: 0.5, z: 0)

[2158] /// A point that's centered vertically on the trailing-front edge of a

[2159] /// view.

[2160] ///

[2161] /// The trailing edge appears on the right in a left-to-right language

[2162] /// environment and on the left in a right-to-left environment.

[2163] public static let trailingFront = UnitPoint3D(x: 1, y: 0.5, z: 1)

[2164] /// A point that's centered horizontally and vertically

[2165] /// on the front face of a view.

[2166] public static let front = UnitPoint3D(x: 0.5, y: 0.5, z: 1)

[2167] /// A point that's centered horizontally and vertically

[2168] /// on the back face of a view.

[2169] public static let back = UnitPoint3D(x: 0.5, y: 0.5, z: 0)

[2170] }

[2171] ``

[2172] As noted earlier, the deprecation of existing `rotation3DEffect` modifiers on mixed reality operating system, along with an mixed reality operating system-exclusive `perspectiveRotationEffect` modifier, produces an inconsistency with other platforms. While we believe this is acceptable for now to help limit the scope of this proposal, we do believe it should be remedied in the future by standardizing on modifiers that are available on all

platforms where possible. While more thinking is needed here, this is one possible solution: Introduce the ``perspectiveRotationEffect`` API and the new ``rotation3DEffect`` overloads that do not take a ``perspective`` value on **all** platforms. Note that these modifiers as proposed use the new ``UnitPoint3D`` type, which necessitates it being made available on all platforms as well. Deprecate the ``rotation3DEffect`` APIs that take a ``perspective`` value on all platforms in favor of using ``perspectiveRotationEffect``. On 2D platforms, we would likely use soft deprecation, as using these APIs in those contexts does not meet the bar of being "actively harmful", while we do believe that to be the case on a 3D platform like mixed reality operating system.

[2173] Animating 3D rotations can be tricky to implement correctly. Typically, 3D rotations are best animated using a spherically linear interpolation, or `[`slerp`]`(<https://en.wikipedia.org/wiki/Slerp>), producing a smooth path between two 3D rotations. This requires being able to interpolate between two values in a way that is possible but difficult to implement using two-dimensional framework's animation system. The ``rotation3DEffect`` modifier can implement this animation behavior automatically for most apps; however, for apps that wish to use similar animations themselves with custom 3D transforms, it may be desirable to provide this more directly. It may be possible to have the ``Rotation3D`` type from the Spatial library conform to ``Animatable`` and implement this behavior, but this has proven difficult to get right in previous attempts. A new ``Animatable`` type may be required instead. This direction needs more investigation; for now, we are using a custom opaque ``Animatable`` type for implementing ``rotation3DEffect``'s ``Animatable`` conformance.

[2174] Two-dimensional framework provides a ``GeometryEffect`` protocol, a kind of ``ViewModifier`` that can be used to apply any 3x3 ``ProjectionTransform`` as a function of the view's 2D size. All geometry effects in two-dimensional framework are currently implemented using that protocol. Behind the scenes, two-dimensional framework on mixed reality operating system also has a ``GeometryEffect3D`` protocol that follows the same design and is how the proposed modifiers like ``offset(z:)`` and ``transform3DEffect`` are implemented, with two key differences: The effect returns a 4x3 ``AffineTransform3D`` instead of a ``ProjectionTransform``. The effect is provided a 3D size to allow it to provide transforms that are a function of the view's depth (e.g., a 3D rotation applied about the view's center in all axes).

[2175] While we considered exposing this protocol publicly as well, we believe this is not necessary with the new `.visualEffect3D` modifier. With this modifier, apps are given even more power than with `GeometryEffect3D`, as they can not only determine the size of a view with the provided `GeometryProxy3D`, but also compute its relative frame or bounds in an ancestor coordinate space like a `ScrollView`.

[2176] In some examples, the feature not supported by using `VisualEffect`'s at this time is the `ignoredByLayout()` modifier on `GeometryEffect` to have a transform effect not affect the view's layout. However, this appears to be quite uncommon to use in most apps, and it could always be added to the `VisualEffect` API later if it proves desirable. While we could also consider exposing `GeometryEffect3D` publicly in the future, it seemed better to encourage use of `.visualEffect3D` for now instead.

[2177] Developers can create standard window and volumetric window experiences on mixed reality operating system and N301, but those windows have three-dimensional bounds defined, which visually clip all of the window view hierarchy's contents. Those constraints however prevent developers from creating rich, immersive experiences, in which three-dimensional entities need to be positioned freely anywhere in the user's environment. Therefore, there is a need for a new high-level view hierarchy container primitive, which developers can define as part of their App, and which provides a space for displaying a view hierarchy without being visually clipped. Such container needs to be able to host an arbitrary two-dimensional framework view hierarchy, including View (proposal), and to be configurable to suit the developer's needs. One requirement of presenting such unbounded container is that all other applications need to be hidden by the system in order to, first and foremost, reduce glitches that can have a negative impact on the user experience due to visual intersections of contents put into the container and those of other applications, and secondly, to give the developers as much space as possible for creating rich and immersive experiences. In line with this, another peculiarity of these containers is that there can always only be none, or one, presented at a time, for the same reasons. Similarly to WindowGroup, clients need to be able to request the presentation of these containers, and to dismiss them, via newly introduced callables. This new container concept will only be available on mixed reality operating system, since there is no need for displaying unbounded content, while running exclusively, on any of our other existing platforms.

[2178] Introduce ImmersiveSpace, a new two-dimensional framework scene type that is aligned with existing window-oriented scene types, but provides the capability to display its contents with no visual border applied. This new type of scene can be used as both a primary scene, or composed with other scene types in an App. Therefore, an ImmersiveSpace can even define the entire interface of an App. Multiple ImmersiveSpace scenes can be defined by the developer, and, similarly to WindowGroup, can be defined using an identifier or a type, that they are capable of handling. The content displayed by an immersive space is defined by a newly introduced ImmersiveSpaceContent protocol, together with a new result builder for it, ImmersiveSpaceContentBuilder . Furthermore, in order to allow View based content, ImmersiveSpaceViewContent is provided, and a convenience initializers for ImmersiveSpace to display View based view hierarchies in ImmersiveSpace. Also introduce new Environment callables, openImmersiveSpace , for presenting a space based on an identifier or a type, and dismissImmersiveSpace , for dismissing the currently opened space. In some examples, to configure the individual characteristics of a space, introduce a new scene modifier that applies to ImmersiveSpace scenes though.

[2179] Similarly to a WindowGroup, developers define one or more immersive spaces in their App:

[2180] ```swift

[2181] @main

[2182] struct MyImmersiveApp: App {

[2183] var body: some Scene {

[2184] ImmersiveSpace {

[2185] ContentView()

[2186] }

[2187] }

[2188] }

[2189] ```

[2190] And, also similarly to WindowGroup , an ImmersiveSpace can be defined with a Hashable & Codable type that then allows developers to present an ImmersiveSpace with a value of that type:

[2191] ```swift

[2192] struct SolarSystem: Codable, Hashable, Identifiable {

[2193] ...

[2194] }

[2195]

[2196] @main

[2197] struct MyImmersiveApp: App {

[2198] var body: some Scene {

[2199] ImmersiveSpace(for: SolarSystem.ID.self) { \$solarSystemID in

[2200] SolarSystemView(solarSystem: \$solarSystem)

[2201] }

[2202] }

[2203] }

[2204] ```

[2205] Immersive spaces can also be defined with a String identifier, in order for developers to disambiguate scenes and be able to open a specific immersive space targeted via an identifier:

[2206] ```swift

[2207] @main

[2208] struct MyImmersiveApp: App {

```

[2209]     var body: some Scene {
[2210]         ImmersiveSpace(id: "solarSystem") {
[2211]             SolarSystemView()
[2212]         }
[2213]     }
[2214] }
[2215] ```

```

[2216] An ImmersiveSpace can be configured using ImmersiveSpace specific modifiers. The immersionStyle modifier allows developers to define which styles the ImmersiveSpace supports. The style has an effect on the appearance of the content of the immersive space. The space starts with a certain style initially, and once the space is open, the style can be changed either programmatically, or by the user's interaction with software or hardware affordances. Both ways of changing the style are constrained by the set of ImmersiveSpace styles assigned to the space. We propose three different styles initially: An immersive space that is presented using the Mixed Immersion style renders its contents on top of the video pass-through of the device. This allows a mixed reality experience, in which the users can still see their surroundings, while also seeing and interacting with virtual objects and user interfaces at the same time. When an immersive space using the Full Immersion style is presented by a developer, all of the video pass-through displaying the physical real world is obscured by the system, except the hands of the user, if configured accordingly (configuration is beyond this proposal). The app does not have to, but is expected to fully cover the entire field of view with the content of the space. This style is meant to be used for fully immersive experiences with the goal of visually isolating the user from the real world. An immersive space using the Progressive Immersion style renders its contents similarly to the Full Immersion style, but does so using a portal effect. The portal offers an angular, more limited view on the virtual contents, while keeping the rest of the physical environment around the portal fully visible. The edges of the portal are feathered. The user can control the level of the immersion via hardware interaction from the initial level up to a level that matches the Full Immersion style. In this version of the platform, there is no way for the

developer of knowing the current immersion level. Apart from the three styles listed above, we also propose an automatic style, `AutomaticImmersionStyle`, with which, once applied, developers get the default immersion style, which resolves as the Mixed Immersion style. The styles which the immersive space allows can be assigned using the new `immersionStyle(selection:,in:)` modifier as follows. It takes a binding for a style, and a variadic list of `ImmersionStyle` instances. The binding can be passed to inner views to both change and observe the currently applied style

[2217] ````swift`

[2218] `@main`

[2219] `struct MyImmersiveApp: App {`

[2220]

[2221] `@State private var currentStyle: ImmersionStyle = .mixed`

[2222]

[2223] `var body: some Scene {`

[2224] `ImmersiveSpace(id: "solarSystem") {`

[2225] `SolarSystemView(immersionStyle: $currentStyle)`

[2226] `}`

[2227] `.immersionStyle(selection: $currentStyle, in: .mixed, .full)`

[2228] `}`

[2229] `}`

[2230] `````

[2231] By default, if no set of allowed styles was assigned, an immersive space will use the Passthrough style as the an allowed one. The variadic list of styles can be empty, which can reduce the call of the modifier to:

[2232] ` ``

[2233] ImmersiveSpace(id: "solarSystem") {

[2234] SolarSystemView(immersionStyle: \$currentStyle)

[2235] }

[2236] .immersionStyle(selection: \$currentStyle)

[2237] ` ``

[2238] Defining an empty list of styles is treated as if a single item list containing the default style, AutomaticImmersionStyle , is provided. If the binding passed to the immersionStyle modifier refers to a style that is not defined in the variadic list of styles, an error is logged, and the first style provided in the list of styles is used instead. Other than that, the order of styles in the list has no meaning. Listing the same style multiple times behaves as listing it once.

[2239] If an immersive space is the scene defined in the App , the space will be presented by the system upon app launch as the primary scene. This requires that the preferred default scene session role specified in the application's Info.plist refers to an immersive space scene session role (see [rdar://89832690](#)). The system will then create a scene that is configured as an immersive space at app launch time, and map it to the first ImmersiveSpace scene defined by the developer that matches. What happens if no such scene can be determined is not part of this proposal. For opening a specific ImmersiveSpace instance either by type or by identifier post-launch, we propose a new openImmersiveSpace callable that is part of the Environment. The different variants of this callable match the ones of the existing openWindow callable. However, as opposed to the window counter part, the openImmersiveSpace callable is designed to work asynchronously (marked as async) in order to allow developers to react to the system having processed the request to open a space, and to return a OpenImmersiveSpaceAction.Result in order to inform the developer about the outcome of the request. Developers can open an immersive space as follows:

[2240] `` swift

[2241] @main

```
[2242]    struct MyImmersiveApp: App {  
[2243]  
[2244]        @Environment(\.openImmersiveSpace) private var openImmersiveSpace  
[2245]  
[2246]        private var solarSystem = SolarSystem(name: "Our Solar System",  
[2247]                                                numberOfPlanets: 8)  
[2248]  
[2249]        var body: some Scene {  
[2250]            WindowGroup {  
[2251]                VStack {  
[2252]                    Button("Show Solar System") {  
[2253]                        Task {  
[2254]                            let result = await openImmersiveSpace(id: "solarSystem")  
[2255]                            if case .error = result {  
[2256]                                print("An error occurred")  
[2257]                            }  
[2258]                        }  
[2259]                    }  
[2260]                    Button("Show Solar System") {  
[2261]                        Task {  
[2262]                            let result = await openImmersiveSpace(value: solarSystem.ID)
```

```

[2263]             if case .error = result {
[2264]                 print("An error occurred")
[2265]             }
[2266]         }
[2267]     }
[2268] }
[2269]
[2270] ImmersiveSpace(id: "solarSystem") {
[2271]     SolarSystemView()
[2272] }
[2273]
[2274] ImmersiveSpace(for: SolarSystem.ID.self) { $solarSystemID in
[2275]     SolarSystemView(solarSystem: $solarSystem)
[2276] }
[2277] }
[2278] }
[2279] ```

```

[2280] One peculiarity of immersive spaces that differentiate them from other Scene types is that one immersive space at a time can be presented. If a developer requests to open an immersive space while one is already shown, or about to be shown, the call to `openImmersiveSpace` will not succeed, an error will be logged, and an error will be return as part of the result of the `openImmersiveSpace` call. Developers are expected to dismiss any currently open immersive space first before opening another one. A request to open an

immersive space can also fail at any time due to the current state of the system. In this case, similarly, an error will be logged, and the result of the call can be used to better understand the reason for the failure.

[2281] As done for window scenes, we are also introducing a dedicated callable for dismissing an immersive space. Similarly to the new `openImmersiveSpace`, the `dismissImmersiveSpace` is offered via the `Environment`, and is designed to work asynchronously in order to be able to react once dismissing is complete. As opposed to `openImmersiveSpace` though, `dismissImmersiveSpace` does not take any arguments, since it always considers the currently opened space as its target, and also does not report any success or error back, as dismissing a space is almost certainly going to succeed.

[2282] ````swift`

[2283] `@main`

[2284] `struct MyImmersiveApp: App {`

[2285] `var body: some Scene {`

[2286] `ImmersiveSpace(id: "solarSystem") {`

[2287] `SolarSystemView()`

[2288] `}`

[2289] `}`

[2290] `}`

[2291]

[2292] `struct SolarSystemView: View {`

[2293]

[2294] `@Environment(\.dismissImmersiveSpace) private var dismissImmersiveSpace`

[2295]

```

[2296]     var body: some View {
[2297]         Button("Dismiss Immersive Space") {
[2298]             Task {
[2299]                 await dismissImmersiveSpace()
[2300]                 print("Dismissing complete")
[2301]             }
[2302]         }
[2303]     }
[2304] }
[2305] ```

```

[2306] Since ImmersiveSpace is a Scene type, an immersive space receives scene phase changes the way other scene phase compatible scenes do as well. The orientation of the coordinate system of an ImmersiveSpace instance, and possible ways to determine or alter it, is not part of this proposal. While arbitrary two-dimensional framework view hierarchies can be used as the content of an immersive space, the use of View as the space's content will be encouraged to developers, as the combined capabilities of the ImmersiveSpace and View act as a great foundation for building immersive experiences on our platform. With View in a space, developers can leverage anchoring and the displaying of unbounded contents while the application is running visually exclusively. The position of a newly opened immersive space is defined by the system and can not be controlled by the application. Developers can expect a space, that is, the origin of its coordinate system, to be positioned at or near what the system has determined as the user's feet's location. If the system considers it to be required, a space can also be positioned slightly moved away from the user's feet. APIs to determine such offset are not part of this proposal. Once an immersive space is presented, the position of it can not be changed by the application. However, the system may re-position at space at any time, if it considered to be valuable to the user, e.g., after the user has physically moved away from the spawning location of the space by a significant amount, or if the user decided to re-

center the experience in front of their current location. APIs to be notified about such system- or user-controlled repositioning events are not part of this proposal.

[2307] When an immersive space is requested to be opened by an application, and the space is presented, the system hides all other applications that may have previously been visually running concurrently. This ensures that the contents of the space do not visually collide or intersect with contents of any other application. It also provides a maximum of visual space for developers for creating immersive experiences. We refer to this as the application running in Exclusive Mode, versus Window Mode for applications that can run visually run side-by-side. Other scenes of the same application remain open though. Developers need to dismiss manually any of their other scenes that should not be visible while an immersive space is opened. The system provides exit actions for users to rapidly switch from Exclusive Mode to Window Mode, in order to have a guaranteed way of returning to unobfuscated video pass-through. This means that invoking such exit actions dismisses any currently presented immersive space. If any other non- ImmersiveSpace kind of scenes are currently visible when invoking, those scenes will remain unchanged. The application will therefore not be dismissed right away, but only following an intentional, second exit action on the hardware button, if a space was opened. Developers can be notified via scene phase changes of the space, if it is dismissed by the user.

[2308] While an immersive space is currently opened, it may obfuscate parts or all of video pass-through, causing the user to be visually disconnected from their physical environment. For safety reasons, and for being able to communicate with people near-by that are stepping into the user's visual field, the system may at any point partially hide the contents, or fully dismiss a space. The system at a later point will restore the full experience, once it is considered safe and appropriate to do so. Ways to be notified via APIs about partial the space contents being hidden by the system temporarily by the system are not part of this proposal. State restoration is a core functionality of the system that ensures, if properly implemented by the developer, a seamless return to the application and the user's recent activity, even in case of application termination. Similarly to other scene types, a space can be restored by the system due to the system's state restoration. This means that if an immersive space was presented when an application was exited due to system events or user interaction, it will still be presented when returning to the application at a later point. If the app was terminated in between, the space will be restored by the system automatically

following the same mechanics as other scene types that support state restoration, and be visible at launch.

[2309] We propose a new ImmersiveSpace type that confirms to the Scene protocol.

[2310] ```swift

[2311]

[2312] /// Conforming types represent items that can be used as the content of an

[2313] /// `ImmersiveSpace`.

[2314] @available(phoneOS, unavailable)

[2315] @available(computerOS, unavailable)

[2316] @available(mediaOS, unavailable)

[2317] @available(wearableOS, unavailable)

[2318] @available(mixed reality operating system 1.0, *)

[2319] public protocol ImmersiveSpaceContent {

[2320] associatedtype Body: ImmersiveSpaceContent

[2321]

[2322] @ImmersiveSpaceContentBuilder public var body: Body { get }

[2323] }

[2324]

[2325] /// A result builder for composing a collection of `ImmersiveSpace` content

[2326] /// items into a single `ImmersiveSpace` content.

[2327] @available(phoneOS, unavailable)

[2328] @available(computerOS, unavailable)

```
[2329]    @available(mediaOS, unavailable)

[2330]    @available(wearableOS, unavailable)

[2331]    @available(mixed reality operating system 1.0, *)

[2332]    @resultBuilder

[2333]    public struct ImmersiveSpaceContentBuilder {

[2334]        public static func buildBlock<Content: ImmersiveSpaceContent>(<
[2335]            _ content: Content
[2336]        ) -> ImmersiveSpaceContent

[2337]    }

[2338]

[2339]    /// A type of `ImmersiveSpace` content that uses a two-dimensional framework
view hierarchy as the

[2340]    /// content. You don't create this type directly. Two-dimensional framework
creates it when you

[2341]    /// construct an `ImmersiveSpace` with view-based content.

[2342]    @available(phoneOS, unavailable)

[2343]    @available(computerOS, unavailable)

[2344]    @available(mediaOS, unavailable)

[2345]    @available(wearableOS, unavailable)

[2346]    @available(mixed reality operating system 1.0, *)

[2347]    public struct ImmersiveSpaceViewContent<Content: View>:
ImmersiveSpaceContent {
```

```
[2348]    }

[2349]

[2350]    /// A scene that presents its content in an unbounded space.

[2351]    ///

[2352]    /// Use an `ImmersiveSpace` as a container for a view hierarchy presented by

[2353]    /// your app. The hierarchy that you declare as the immersive space's content

[2354]    /// serves as a template for it:

[2355]    ///

[2356]    /// @main

[2357]    /// struct SolarSystemApp: App {

[2358]    ///     var body: some Scene {

[2359]    ///         ImmersiveSpace {

[2360]    ///             SolarSystem()

[2361]    ///         }

[2362]    ///     }

[2363]    /// }

[2364]    ///

[2365]    /// When your app opens an immersive space, the system hides all other visible

[2366]    /// apps. two-dimensional framework allows one opened ``ImmersiveSpace`` at a

    time.

[2367]    /// Close any currently opened immersive space before opening another one.

[2368]    ///
```

- [2369] /// You can use ``ImmersiveSpace`` as the main interface of an app.
- [2370] /// For bounded scenes, use ```WindowGroup``` instead.
- [2371] /// For document-based apps, use a ```DocumentGroup``` instead.
- [2372] ///
- [2373] /// You can optionally define an ``ImmersiveSpace`` to present data of a type that
- [2374] /// conforms to both
<doc://com.apple.documentation/documentation/Swift/Hashable>
- [2375] /// and <doc://com.apple.documentation/documentation/Swift/Codable>.
- [2376] ///
- [2377] /// If you pass a value to the ```EnvironmentValues/openImmersiveSpace``` action
- [2378] /// to open an immersive space, the space receives a binding to the presented
- [2379] /// value as its root view.
- [2380] /// The system persists the value of the binding for the purposes of state
- [2381] /// restoration. When restoring the space, the system decodes the value and uses
- [2382] /// it to set the binding. If decoding fails, the system sets the binding to
- [2383] /// either the default value or ``nil``.
- [2384] ///
- [2385] /// Prefer using lightweight data for your presentation value.
- [2386] /// For structured model values that conform to
- [2387] /// <doc://com.apple.documentation/documentation/Swift/Identifiable>,
- [2388] /// the value's identifier works well.
- [2389] /// For example, you can define a button that opens the specified solar system

```
[2390]    /// value in a new immersive space:

[2391]    ///

[2392]    ///    @main

[2393]    ///    struct SolarSystemApp: App {

[2394]    ///        var body: some Scene {

[2395]    ///            ImmersiveSpace(for: SolarSystem.ID.self) { $solarSystemID in

[2396]    ///                // ...

[2397]    ///            }

[2398]    ///        }

[2399]    ///    }

[2400]    ///

[2401]    ///    struct NewSolarSystemImmersiveSpace: View {

[2402]    ///        var solarSystem: SolarSystem

[2403]    ///        @Environment(\.openImmersiveSpace) private var openImmersiveSpace

[2404]    ///

[2405]    ///        var body: some View {

[2406]    ///            Button("Present Solar System In New Immersive Space") {

[2407]    ///                Task {

[2408]    ///                    await openImmersiveSpace(value: solarSystem.ID)

[2409]    ///                }

[2410]    ///            }
```

```
[2411]    ///    }
[2412]    ///    }
[2413]    ///
[2414]    @available(phoneOS, unavailable)
[2415]    @available(computerOS, unavailable)
[2416]    @available(mediaOS, unavailable)
[2417]    @available(wearableOS, unavailable)
[2418]    @available(mixed reality operating system 1.0, *)
[2419]    public struct ImmersiveSpace<
[2420]        Content: ImmersiveSpaceContent, Data: Hashable & Codable
[2421]    >: Scene {
[2422]
[2423]        /// Creates an immersive space.
[2424]        ///
[2425]        /// The space uses the specified content builder to form the content.
[2426]        ///
[2427]        /// - Parameters:
[2428]        ///   - content: An immersive space content builder that defines the content
[2429]        ///   of the space.
[2430]        public init(@ImmersiveSpaceContentBuilder content: () -> Content)
[2431]        where Data == Never
```

[2432]

[2433] /// Creates the immersive space associated with the specified identifier.

[2434] ///

[2435] /// The space uses the specified content builder to form the content.

[2436] ///

[2437] /// - Parameters:

[2438] /// - id: A string that uniquely identifies the immersive space.

[2439] /// Identifiers must be unique among the immersive spaces in your app.

[2440] /// - content: An immersive space content builder that defines the content

[2441] /// of the space.

[2442] public init(
[2443] id: String,
[2444] @ImmersiveSpaceContentBuilder content: () -> Content
[2445]) where Data == Never

[2446]

[2447] /// Creates the immersive space associated with an identifier for a

[2448] /// specified type of presented data.

[2449] ///

[2450] /// The space uses the specified content builder to form the content.

[2451] ///

[2452] /// Your app invokes this initializer when it presents a value

[2453] /// of the specified `type` using the

[2454] /// ``EnvironmentValues/openImmersiveSpace`` action.

[2455] ///

[2456] /// - Parameters:

[2457] /// - id: A string that uniquely identifies the immersive space.

[2458] /// Identifiers must be unique among the immersive spaces in your app.

[2459] /// - type: The type of presented data this space accepts.

[2460] /// - content: An immersive space content builder that creates the content

[2461] /// for each instance of the space. The closure receives a binding to

[2462] /// the value passed via ``EnvironmentValues/openImmersiveSpace``.

[2463] /// The system automatically persists and restores the value of this

[2464] /// binding during state restoration.

[2465] public init(
[2466] id: String, for type: Data.Type,
[2467] @ImmersiveSpaceContentBuilder content: @escaping (Binding<Data?>)
[2468] -> Content
[2469])
[2470]
[2471] /// Creates the immersive space for a specified type of presented data.
[2472] ///
[2473] /// The space uses the specified content builder to form the content.

[2474] ///

[2475] /// Your app invokes this initializer when it presents a value

[2476] /// of the specified `type` using the

[2477] /// ``EnvironmentValues/openImmersiveSpace``

[2478] /// action.

[2479] ///

[2480] /// - Parameters:

[2481] /// - type: The type of presented data this space accepts.

[2482] /// - content: An immersive space content builder that defines the content

[2483] /// of the space. The closure receives a binding to the value passed via

[2484] /// ``EnvironmentValues/openImmersiveSpace``. The value of this binding

[2485] /// will be persisted and restored during state restoration of the

[2486] /// space.

[2487] public init(
[2488] for type: Data.Type,
[2489] @ImmersiveSpaceContentBuilder content: @escaping (Binding<Data?>)
[2490] -> Content
[2491])
[2492]

[2493] /// Creates the immersive space associated with an identifier for a

[2494] /// specified type of presented data, and a default value, if the data is

[2495] /// not set.

[2496] ///

[2497] /// The space uses the specified content builder to form the content.

[2498] ///

[2499] /// Your app invokes this initializer when it presents a value

[2500] /// of the specified `type` using the

[2501] /// ``EnvironmentValues/openImmersiveSpace`` action.

[2502] ///

[2503] /// - Parameters:

[2504] /// - id: A string that uniquely identifies the immersive space.

[2505] /// Identifiers must be unique among the immersive spaces in your app.

[2506] /// - type: The type of presented data this space accepts.

[2507] /// - content: An immersive space content builder that defines the content

[2508] /// of the space. The closure receives a binding to the value passed via

[2509] /// ``EnvironmentValues/openImmersiveSpace``. The value of this binding

[2510] /// will be persisted and restored during state restoration of the

[2511] /// space.

[2512] /// - defaultValue: A closure that returns a default value to present,

[2513] /// which will be called in situations where the immersive space is

[2514] /// presented without a value provided.

[2515] public init(

```
[2516]      id: String, for type: Data.Type = Data.self,
[2517]      @ImmersiveSpaceContentBuilder content: @escaping (Binding<Data>)
[2518]          -> Content,
[2519]      defaultValue: @escaping () -> Data
[2520]  )
[2521]
[2522]  /// Creates an immersive space.
[2523]  ///
[2524]  /// The immersive space uses the specified content builder to form the
[2525]  /// content.
[2526]  ///
[2527]  /// Your app invokes this initializer when it presents a value
[2528]  /// of the specified `type` using the
[2529]  /// ``EnvironmentValues/openImmersiveSpace`` action.
[2530]  ///
[2531]  /// - Parameters:
[2532]  ///   - type: The type of presented data this space accepts.
[2533]  ///   - content: An immersive space content builder that defines the content
[2534]  ///     of the space. The closure receives a binding to the value passed via
[2535]  ///     ``EnvironmentValues/openImmersiveSpace``. The value of this binding
[2536]  ///     will be persisted and restored during state restoration of the
```

```

[2537]    ///    space.

[2538]    ///    - defaultValue: A closure that returns a default value to present,

[2539]    ///    which will be called in situations where the immersive space is

[2540]    ///    presented without a value provided.

[2541]    public init(

[2542]        for type: Data.Type = Data.self,

[2543]        @ImmersiveSpaceContentBuilder content: @escaping (Binding<Data>)

[2544]            -> Content,

[2545]        defaultValue: @escaping () -> Data

[2546]    )

[2547]    }

[2548]

[2549]    extension ImmersiveSpace where Content == ImmersiveSpaceViewContent<V> {

[2550]

[2551]        /// A convenience for creating an immersive space using view-based content.

[2552]        ///

[2553]        /// The immersive space uses the specified view as a template to form the

[2554]        /// content of the space.

[2555]        ///

[2556]        /// - Parameters:

[2557]        ///    - content: A view builder that defines the content for each

```

[2558] /// instance of the immersive space.

[2559] public init<V: View>(
[2560] @ViewBuilder content: () -> V
[2561]) where Data == Never
[2562]

[2563] /// A convenience for creating an immersive space using view-based content.

[2564] /// Creates the immersive space associated with the specified identifier.

[2565] ///

[2566] /// The immersive space uses the specified view as a template to form the
[2567] /// content of the space.
[2568] ///

[2569] /// - Parameters:

[2570] /// - id: A string that uniquely identifies the immersive space.

[2571] /// Identifiers must be unique among the immersive spaces in your app.

[2572] /// - content: A view builder that defines the content for each
[2573] /// instance of the immersive space.

[2574] public init<V: View>(
[2575] id: String,
[2576] @ViewBuilder content: () -> V
[2577]) where Data == Never
[2578]

[2579] /// A convenience for creating an immersive space using view-based content.

[2580] /// Creates the immersive space associated with an identifier for a

[2581] /// specified type of presented data.

[2582] ///

[2583] /// The immersive space uses the specified view as a template to form the

[2584] /// content of the space.

[2585] ///

[2586] /// Your app invokes this initializer when it presents a value

[2587] /// of the specified `type` using the

[2588] /// ``EnvironmentValues/openImmersiveSpace`` action.

[2589] ///

[2590] /// - Parameters:

[2591] /// - id: A string that uniquely identifies the immersive space.

[2592] /// Identifiers must be unique among the immersive spaces in your app.

[2593] /// - type: The type of presented data this immersive space accepts.

[2594] /// - content: A view builder that defines the content for each

[2595] /// instance of the immersive space. The closure receives a binding to

[2596] /// the value passed via ``EnvironmentValues/openImmersiveSpace``.

[2597] /// The system automatically persists and restores the value of this

[2598] /// binding during state restoration.

[2599] public init<V: View>(

[2600] id: String, for type: Data.Type,

[2601] @ViewBuilder content: @escaping (Binding<Data?>) -> V

[2602])

[2603]

[2604] /// A convenience for creating an immersive space using view-based content.

[2605] /// Creates the immersive space for a specified type of presented data.

[2606] ///

[2607] /// The immersive space uses the specified view as a template to form the

[2608] /// content of the space.

[2609] ///

[2610] /// Your app invokes this initializer when it presents a value

[2611] /// of the specified `type` using the

[2612] /// ``EnvironmentValues/openImmersiveSpace`` action.

[2613] ///

[2614] /// - Parameters:

[2615] /// - type: The type of presented data this immersive space accepts.

[2616] /// - content: A view builder that defines the content for each

[2617] /// instance of the immersive space. The closure receives a binding to

[2618] /// the value passed via ``EnvironmentValues/openImmersiveSpace``.

[2619] /// The system automatically persists and restores the value of this

[2620] /// binding during state restoration.

```
[2621]    public init<V: View>(  
[2622]        for type: Data.Type,  
[2623]        @ViewBuilder content: @escaping (Binding<Data?>) -> V  
[2624]    )  
[2625]  
[2626]    /// A convenience for creating an immersive space using view-based content.  
[2627]    /// Creates the immersive space associated with an identifier for a  
[2628]    /// specified type of presented data, and a default value, if the data is  
[2629]    /// not set.  
[2630]    ///  
[2631]    /// The immersive space uses the specified view as a template to form the  
[2632]    /// content of the space.  
[2633]    ///  
[2634]    /// Your app invokes this initializer when it presents a value  
[2635]    /// of the specified `type` using the  
[2636]    /// ``EnvironmentValues/openImmersiveSpace`` action.  
[2637]    ///  
[2638]    /// - Parameters:  
[2639]    ///   - id: A string that uniquely identifies the immersive space.  
[2640]    ///   Identifiers must be unique among the immersive spaces in your app.  
[2641]    ///   - type: The type of presented data this immersive space accepts.
```

[2642] /// - content: A view builder that defines the content for each

[2643] /// instance of the immersive space. The closure receives a binding to

[2644] /// the value passed via ``EnvironmentValues/openImmersiveSpace``.

[2645] /// The system automatically persists and restores the value of this

[2646] /// binding during state restoration.

[2647] /// - defaultValue: A closure that returns a default value to present,

[2648] /// which will be called in situations where the immersive space is

[2649] /// presented without a value provided.

[2650] public init<V: View>(
[2651] id: String, for type: Data.Type = Data.self,
[2652] @ViewBuilder content: @escaping (Binding<Data>) -> V,
[2653] defaultValue: @escaping () -> Data
[2654])
[2655]
[2656] /// A convenience for creating an immersive space using view-based content.
[2657] /// Creates an immersive space for a specified type of presented data, and a
[2658] /// default value, if the data is not set.
[2659] ///
[2660] /// The immersive space uses the specified view as a template to form the
[2661] /// content of the space.
[2662] ///

```

[2663]    /// Your app invokes this initializer when it presents a value
[2664]    /// of the specified `type` using the
[2665]    /// ``EnvironmentValues/openImmersiveSpace`` action.
[2666]    ///
[2667]    /// - Parameters:
[2668]    ///   - type: The type of presented data this immersive space accepts.
[2669]    ///   - content: A view builder that defines the content for each
[2670]    ///     instance of the immersive space. The closure receives a binding to
[2671]    ///     the value passed via ``EnvironmentValues/openImmersiveSpace``.
[2672]    ///   The system automatically persists and restores the value of this
[2673]    ///   binding during state restoration.
[2674]    ///   - defaultValue: A closure that returns a default value to present,
[2675]    ///     which will be called in situations where the immersive space is
[2676]    ///     presented without a value provided.
[2677]    public init<V: View>(
[2678]        for type: Data.Type = Data.self,
[2679]        @ViewBuilder content: @escaping (Binding<Data>) -> V,
[2680]        defaultValue: @escaping () -> Data
[2681]    )
[2682]    }
[2683]    ```

```

[2684] In order to define the allowed styles of an ImmersiveSpace , we propose ImmersionStyle and a new Scene modifier. This modifier acts as a no-op on scenes that are not an ImmersiveSpace:

[2685] ```swift

[2686] /// The allowed styles of an `ImmersiveSpace`.

[2687] ///

[2688] /// Apply one or more styles that conform to this protocol to an

[2689] /// `ImmersiveSpace` instance to configure the appearance and behavior of the

[2690] /// space using the ``Scene/immersionStyle(selection:,in:)`` scene modifier.

[2691]

[2692] /// A type that applies a style to an `ImmersiveSpace`.

[2693] /// The style can change the appearance and behavior of the immersive space.

[2694] ///

[2695] /// To configure the current style of an immersive space, use the

[2696] /// ``Scene/immersionStyle(selection:,in:)`` modifier. Specify a style that

[2697] /// conforms to `ImmersionStyle` when creating a Space.

[2698] @available(phoneOS, unavailable)

[2699] @available(computerOS, unavailable)

[2700] @available(mediaOS, unavailable)

[2701] @available(wearableOS, unavailable)

[2702] @available(mixed reality operating system 1.0, *)

[2703] public protocol ImmersionStyle { }

[2704]

[2705] @available(phoneOS, unavailable)

[2706] @available(computerOS, unavailable)

[2707] @available(mediaOS, unavailable)

[2708] @available(wearableOS, unavailable)

[2709] @available(mixed reality operating system 1.0, *)

[2710] extension Scene {

[2711] /// Sets the allowed styles for the immersive space.

[2712] /// - Parameters:

[2713] /// - selection: A binding to the effective style used by the space.

[2714] /// - styles: The list of styles that the immersive space allows.

[2715] public func immersionStyle(selection: Binding<any ImmersionStyle>,

[2716] in styles: any ImmersionStyle...) -> some Scene

[2717] }

[2718] ``

[2719] Furthermore, we introduce three styles:

[2720] ``swift

[2721] /// An immersion style that displays unbounded content intermixed with other

[2722] /// app content, along with pass-through video.

[2723] ///

[2724] /// Use ``ImmersionStyle/mixed`` to construct this style.

[2725] @available(phoneOS, unavailable)

[2726] @available(computerOS, unavailable)

[2727] @available(mediaOS, unavailable)

[2728] @available(wearableOS, unavailable)

[2729] @available(mixed reality operating system 1.0, *)

[2730] public struct MixedImmersionStyle: ImmersionStyle {

[2731] public init()

[2732] }

[2733]

[2734] /// An immersion style that displays unbounded content that obscures

[2735] /// pass-through video.

[2736] ///

[2737] /// Your space's content fully obscures the pass-through video

[2738] /// except for the user's hands, if configured accordingly.

[2739] ///

[2740] /// Use ``ImmersionStyle/full`` to construct this style.

[2741] @available(phoneOS, unavailable)

[2742] @available(computerOS, unavailable)

[2743] @available(mediaOS, unavailable)

[2744] @available(wearableOS, unavailable)

[2745] @available(mixed reality operating system 1.0, *)

```
[2746]    public struct FullImmersionStyle: ImmersionStyle {  
[2747]        public init()  
[2748]    }  
[2749]  
[2750]    /// An immersion style in which the content is displayed with no clipping  
[2751]    /// boundaries applied. A radial portal effect is used initially. Users can then  
[2752]    /// interactively dial in and out to switch between the portal style and a style  
[2753]    /// that matches the `FullImmersionStyle`. In the case of the latter  
[2754]    /// pass-through is obscured, except for the user's hands, if configured  
[2755]    /// accordingly.  
[2756]    ///  
[2757]    /// Use ``ImmersionStyle/progressive`` to construct this style.  
[2758]    @available(phoneOS, unavailable)  
[2759]    @available(computerOS, unavailable)  
[2760]    @available(mediaOS, unavailable)  
[2761]    @available(wearableOS, unavailable)  
[2762]    @available(mixed reality operating system 1.0, *)  
[2763]    public struct ProgressiveImmersionStyle: ImmersionStyle {  
[2764]        public init()  
[2765]    }  
[2766]    ``
```

[2767] Apart from the three concrete styles, we introduce the automatic style:

[2768] ```swift

[2769] /// The default style of immersive spaces.

[2770] ///

[2771] /// By default, two-dimensional framework uses the ``ImmersionStyle/mixed``
style as

[2772] /// the automatic style.

[2773] ///

[2774] /// Use ``ImmersionStyle/automatic`` to construct this style.

[2775] @available(phoneOS, unavailable)

[2776] @available(computerOS, unavailable)

[2777] @available(mediaOS, unavailable)

[2778] @available(wearableOS, unavailable)

[2779] @available(mixed reality operating system 1.0, *)

[2780] public struct AutomaticImmersionStyle: ImmersionStyle {

[2781] public init()

[2782] }

[2783] ``

[2784] These styles are also made available via the following extensions:

[2785] ```swift

[2786] @available(phoneOS, unavailable)

[2787] @available(computerOS, unavailable)

[2788] @available(mediaOS, unavailable)

[2789] @available(wearableOS, unavailable)

[2790] @available(mixed reality operating system 1.0, *)

[2791] extension ImmersionStyle where Self == MixedImmersionStyle {

[2792] /// An immersion style that displays unbounded content intermixed with other

[2793] /// app content, along with pass-through video.

[2794] public static var mixed: MixedImmersionStyle

[2795] }

[2796]

[2797] /// An immersion style that displays unbounded content that obscures

[2798] /// pass-through video.

[2799] ///

[2800] /// Your space's content fully obscures the pass-through video

[2801] /// except for the user's hands, if configured accordingly.

[2802] @available(phoneOS, unavailable)

[2803] @available(computerOS, unavailable)

[2804] @available(mediaOS, unavailable)

[2805] @available(wearableOS, unavailable)

[2806] @available(mixed reality operating system 1.0, *)

[2807] extension ImmersionStyle where Self == FullImmersionStyle {

[2808] /// An immersion style that displays unbounded content that obscures

[2809] /// pass-through video, except for the user's hands, if configured

[2810] /// accordingly.

[2811] public static var full: FullImmersionStyle

[2812] }

[2813]

[2814] /// An immersion style that displays unbounded content in a portal.

[2815] @available(phoneOS, unavailable)

[2816] @available(computerOS, unavailable)

[2817] @available(mediaOS, unavailable)

[2818] @available(wearableOS, unavailable)

[2819] @available(mixed reality operating system 1.0, *)

[2820] extension ImmersionStyle where Self == ProgressiveImmersionStyle {

[2821] /// An immersion style in which the content is displayed with no clipping

[2822] /// boundaries applied. A radial portal effect is used initially. Users can

[2823] /// then interactively dial in and out to switch between the portal style

[2824] /// and a style that matches the `FullImmersionStyle`. In the

[2825] /// case of the latter, passthrough is obscured, except for the user's

[2826] /// hands, if configured accordingly.

[2827] public static var progressive: ProgressiveImmersionStyle

[2828] }

[2829]

[2830] /// The default immersion style.

[2831] ///

[2832] /// By default, two-dimensional framework uses the ``ImmersionStyle/mixed`` style as

[2833] /// the automatic style.

[2834] @available(phoneOS, unavailable)

[2835] @available(computerOS, unavailable)

[2836] @available(mediaOS, unavailable)

[2837] @available(wearableOS, unavailable)

[2838] @available(mixed reality operating system 1.0, *)

[2839] extension ImmersionStyle where Self == AutomaticImmersionStyle {

[2840] /// The default immersion style.

[2841] public static var automatic: AutomaticImmersionStyle

[2842] }

[2843] ``

[2844] If a user has not defined any allowed style on an immersive space, the automatic style will be used as the only allowed style, which resolves to the same runtime behavior as the Mixed Immersion style. For opening an immersive space, we propose adding a new `OpenImmersiveSpaceAction`, that is made available as a newly introduced `openImmersiveSpace` callable. For accessing the `openImmersiveSpace` callable, we propose augmenting the `Environment` with a new `openImmersiveSpace` key:

[2845] ``swift

[2846] /// An action that presents an `ImmersiveSpace` instance.

[2847] ///

[2848] /// If the targeted immersive space accepts data, then the system provides the

[2849] /// value that you pass to the action to the space's content closure.

[2850] ///

[2851] /// Use the ``EnvironmentValues/openImmersiveSpace`` environment value to get

[2852] /// the instance of this structure for a given ``Environment``. Then call the

[2853] /// instance to present a space. You call the instance directly because it

[2854] /// defines a ``OpenImmersiveSpaceAction/callAsFunction(id:`` method that
Swift

[2855] /// calls when you call the instance.

[2856] ///

[2857] /// For example, you can define a button that opens the specified solar system

[2858] /// item in a new immersive space:

[2859] ///

[2860] /// @main

[2861] /// struct SolarSystemApp: App {

[2862] /// var body: some Scene {

[2863] /// ImmersiveSpace(for: SolarSystem.ID.self) { \$solarSystemID in

[2864] /// // ...

[2865] /// }

[2866] /// }

[2867] /// }

```
[2868]    ///
[2869]    ///    struct NewSolarSystemImmersiveSpace: View {
[2870]    ///        var solarSystem: SolarSystem
[2871]    ///        @Environment(\.openImmersiveSpace) private var openImmersiveSpace
[2872]    ///
[2873]    ///        var body: some View {
[2874]    ///            Button("Present Solar System In New Immersive Space") {
[2875]    ///                Task {
[2876]    ///                    await openImmersiveSpace(value: solarSystem.ID)
[2877]    ///                }
[2878]    ///            }
[2879]    ///        }
[2880]    ///    }
[2881]    ///
[2882]    /// For best performance, use lightweight data for the presentation value.
[2883]    /// For structured model values that conform to
[2884]    /// <doc://com.apple.documentation/documentation/Swift/Identifiable>,
[2885]    /// the value's identifier makes a good presentation value.
[2886]    @available(phoneOS, unavailable)
[2887]    @available(computerOS, unavailable)
[2888]    @available(mediaOS, unavailable)
```

```
[2889]    @available(wearableOS, unavailable)

[2890]    @available(mixed reality operating system 1.0, *)

[2891]    public struct OpenImmersiveSpaceAction {

[2892]

[2893]        /// Represents the result of opening an ``ImmersiveSpace`` scene.

[2894]        public enum Result: Sendable {

[2895]            /// Opening the immersive space succeeded.

[2896]            case opened

[2897]            /// Opening the immersive space failed since the user cancelled the

[2898]            /// request.

[2899]            case userCancelled

[2900]            /// Opening the immersive space failed since the system cannot fulfill

[2901]            /// the request.

[2902]            case error

[2903]        }

[2904]

[2905]        /// Presents the immersive space that handles the type of the presented

[2906]        /// value.

[2907]        ///

[2908]        /// The call returns either once the immersive space is presented, or once

[2909]        /// an error occurred.
```

[2910] ///

[2911] /// Two-dimensional framework calls it when you

[2912] /// call the ``EnvironmentValues/openImmersiveSpace`` action:

[2913] ///

[2914] /// await openImmersiveSpace(value: solarSystem.ID)

[2915] ///

[2916] /// If an immersive space is already presented, calling this action has no

[2917] /// effect.

[2918] /// Dismiss any currently presented immersive space prior to presenting

[2919] /// another one. Otherwise, the system creates a new immersive space and

[2920] /// passes its root view a binding to the specified value.

[2921] ///

[2922] /// For best performance, use lightweight data for the presentation value.

[2923] /// For structured model values that conform to

[2924] /// <doc://com.apple.documentation/documentation/Swift/Identifiable>,

[2925] /// the value's identifier makes a good presentation value.

[2926] ///

[2927] /// - Parameters:

[2928] /// - value: The value to present.

[2929] @discardableResult

[2930] public func callAsFunction<D: Hashable & Codable>(value: D) async ->
Result

[2931]

[2932] /// Presents an immersive space for the scene with the specified identifier.

[2933] ///

[2934] /// The call returns either once the space is presented, or once an error

[2935] /// occurred.

[2936] ///

[2937] /// If an immersive space is already presented, calling this action has no

[2938] /// effect.

[2939] /// Dismiss any currently presented immersive space prior to presenting

[2940] /// another one.

[2941] ///

[2942] /// If the immersive space presents data, the system provides the default

[2943] /// value or `nil` to the the space's root view.

[2944] ///

[2945] /// - Parameters:

[2946] /// - id: The identifier of the immersive space to present.

[2947] @discardableResult

[2948] public func callAsFunction(id: String) async -> Result

[2949]

[2950] /// Presents the immersive space that your app defines for the specified

[2951] /// identifier and that handles the type of the presented value.

[2952] ///

[2953] /// The call returns either once the space is presented, or once an error

[2954] /// occurred.

[2955] ///

[2956] /// Don't call this method directly. Two-dimensional framework calls it when
you

[2957] /// call the ``EnvironmentValues/openImmersiveSpace`` action:

[2958] ///

[2959] /// await openImmersiveSpace(id: "solarSystem", value: solarSystem.ID)

[2960] ///

[2961] /// If an immersive space is already presented, calling this action has no

[2962] /// effect.

[2963] /// Dismiss any currently presented immersive space prior to presenting

[2964] /// another one.

[2965] /// Otherwise, the system creates a new immersive space and passes its root

[2966] /// view a binding to the specified value.

[2967] ///

[2968] /// For best performance, use lightweight data for the presentation value.

[2969] /// For structured model values that conform to

[2970] /// <doc://com.apple.documentation/documentation/Swift/Identifiable>,

[2971] /// the value's identifier makes a good presentation value.

[2972] ///

[2973] /// - Parameters:

[2974] /// - id: The identifier of the immersive space to present.

[2975] /// - value: The value to present.

[2976] @discardableResult

[2977] public func callAsFunction<D: Hashable & Codable>(
[2978] id: String, value: D
[2979]) async -> Result
[2980] }
[2981] ```\n

[2982] We would then provide this action via the Environment as follows:

[2983] ```swift

[2984] @available(phoneOS, unavailable)

[2985] @available(computerOS, unavailable)

[2986] @available(mediaOS, unavailable)

[2987] @available(wearableOS, unavailable)

[2988] @available(mixed reality operating system 1.0, *)

[2989] extension EnvironmentValues {
[2990] /// An action that presents an immersive space.
[2991] ///
[2992] /// The call returns either once the space is presented, or once an error
[2993] /// occurred.

[2994] ///

[2995] /// If the targeted scene accepts data, then the value passed to the action

[2996] /// that you pass to the action to the scene's content closure.

[2997] ///

[2998] /// Use the ``EnvironmentValues/openImmersiveSpace`` environment value to

[2999] /// get the instance of this structure for a given ``Environment``. Then

[3000] /// call the instance to present a space. You call the instance directly

[3001] /// because it defines a ``OpenImmersiveSpaceAction/callAsFunction(id:)``

[3002] /// method that Swift calls when you call the instance.

[3003] ///

[3004] /// For example, you can define a button that opens the specified solar

[3005] /// system value in a new space.

[3006] ///

[3007] /// @main

[3008] /// struct SolarSystemApp: App {

[3009] /// var body: some Scene {

[3010] /// ImmersiveSpace(for: SolarSystem.ID.self) { \$solarSystemID in

[3011] /// // ...

[3012] /// }

[3013] /// }

[3014] /// }

```
[3015]    ///
[3016]    ///    struct NewSolarSystemImmersiveSpace: View {
[3017]    ///        var solarSystem: SolarSystem
[3018]    ///        @Environment(\.openImmersiveSpace)
[3019]    ///        private var openImmersiveSpace
[3020]    ///
[3021]    ///        var body: some View {
[3022]    ///            Button("Present Solar System In New Immersive Space") {
[3023]    ///                Task {
[3024]    ///                    await openImmersiveSpace(value: solarSystem.ID)
[3025]    ///                }
[3026]    ///            }
[3027]    ///        }
[3028]    ///    }
[3029]    ///
[3030]    public var openImmersiveSpace: OpenImmersiveSpaceAction { get }
[3031]    }
[3032]
[3033]    @available(*, unavailable)
[3034]    extension OpenImmersiveSpaceAction: Sendable {}
[3035]    ``
```

[3036] In order to dismiss an immersive space, we introduce the dismissImmersiveSpace callable.

[3037] ```swift

[3038] /// An action that dismisses an immersive space.

[3039] ///

[3040] /// Use the ``EnvironmentValues/dismissImmersiveSpace`` environment value to get

[3041] /// the instance of this structure for a given ``Environment``. Then call the

[3042] /// instance to dismiss the currently opened space. You call the instance

[3043] /// directly because it defines a

[3044] /// ``DismissImmersiveSpaceAction/callAsFunction()`` method that Swift calls

[3045] /// when you call the instance.

[3046] ///

[3047] /// For example, you can define a button that closes the currently opened

[3048] /// immersive space:

[3049] ///

[3050] /// @main

[3051] /// struct MyApp: App {

[3052] /// var body: some Scene {

[3053] /// WindowGroup {

[3054] /// ContentView()

[3055] /// }

```
[3056]    ///      ImmersiveSpace(id: "solarSystem") {
[3057]    ///          SolarSystemView()
[3058]    ///      }
[3059]    ///  }
[3060]    /// }
[3061]    ///
[3062]    /// struct DismissImmersiveSpaceButton: View {
[3063]    ///     @Environment(\.dismissImmersiveSpace)
[3064]    ///     private var dismissImmersiveSpace
[3065]    ///
[3066]    ///     var body: some View {
[3067]    ///         Button("Close Solar System Immersive Space") {
[3068]    ///             Task {
[3069]    ///                 await dismissImmersiveSpace()
[3070]    ///                 print("Dismissing Immersive Space complete")
[3071]    ///             }
[3072]    ///         }
[3073]    ///     }
[3074]    /// }
[3075]    @available(phoneOS, unavailable)
[3076]    @available(computerOS, unavailable)
```

```
[3077]    @available(mediaOS, unavailable)

[3078]    @available(wearableOS, unavailable)

[3079]    @available(mixed reality operating system 1.0, *)

[3080]    public struct DismissImmersiveSpaceAction {

[3081]        /// Dismisses the currently opened immersive space.

[3082]        ///

[3083]        /// The call returns once the space is dismissed.

[3084]        ///

[3085]        /// Don't call this method directly. Two-dimensional framework calls it when
you
[3086]        /// call the ``EnvironmentValues/dismissImmersiveSpace`` action:

[3087]        ///

[3088]        ///     await dismissImmersiveSpace()

[3089]        ///

[3090]        /// For information about how Swift uses the `callAsFunction()` method to

[3091]        /// simplify call site syntax, see

[3092]        /// [Methods with Special Names](https://docs.swift.org/swift-
book/ReferenceManual/Declarations.html#ID622)

[3093]        /// in *The Swift Programming Language*.

[3094]        public func callAsFunction() async

[3095]    }

[3096]
```

[3097] @available(*, unavailable)

[3098] extension DismissImmersiveSpaceAction: Sendable {}

[3099]

[3100] @available(phoneOS 17.0, *)

[3101] @available(computerOS 14.0, *)

[3102] @available(mediaOS, unavailable)

[3103] @available(wearableOS, unavailable)

[3104] @available(mixed reality operating system 1.0, *)

[3105] extension EnvironmentValues {

[3106] /// An immersive space dismissal action stored in a view's environment.

[3107] ///

[3108] /// The call returns once the immersive space is dismissed.

[3109] ///

[3110] /// Use the `dismissImmersiveSpace` environment value to get an

[3111] /// ``DismissImmersiveSpaceAction`` instance for a given ``Environment``.

[3112] /// Then call the instance to dismiss a space. You call the instance

[3113] /// directly because it defines a

[3114] /// ``DismissImmersiveSpaceAction/callAsFunction()`` method that Swift calls

[3115] /// when you call the instance.

[3116] ///

[3117] /// For example, you can define a button that dismisses an immersive space:

```
[3118]    ///
[3119]    ///  @main
[3120]    ///  struct MyApp: App {
[3121]    ///      var body: some Scene {
[3122]    ///          WindowGroup {
[3123]    ///              ContentView()
[3124]    ///          }
[3125]    ///          ImmersiveSpace(id: "solarSystem") {
[3126]    ///              SolarSystemView()
[3127]    ///          }
[3128]    ///      }
[3129]    ///  }
[3130]    ///
[3131]    ///  struct DismissImmersiveSpaceButton: View {
[3132]    ///      @Environment(\.dismissImmersiveSpace)
[3133]    ///      private var dismissImmersiveSpace
[3134]    ///
[3135]    ///      var body: some View {
[3136]    ///          Button("Close Solar System Immersive Space") {
[3137]    ///              Task {
[3138]    ///                  await dismissImmersiveSpace()
```

```

[3139]      ///          print("Dismissing complete")
[3140]      ///          }
[3141]      ///          }
[3142]      ///      }
[3143]      ///  }
[3144]      public var dismissImmersiveSpace: DismissImmersiveSpaceAction { get }
[3145]  }
[3146]  ```

```

[3147] Similarly to dismissWindow, developers can use the dismissBehavior transaction property that allows for dismissing an immersive space regardless of the presence of modals.

[3148] It may be an option to loosen the current requirement of the application entering Exclusive Mode when an ImmersiveSpace is presented. This could be useful for applications where the space has a rather decorative character, and other applications running visually side-by-side at the same time may be a valid request. In certain use cases, it may be desirable to have control over the positioning and orientation of the space when opening it, and when the user physically, or any of the other scenes, are being moved around. For example, a space that displays simply visual effects, such as glow, fire or fog, may require a positioning system in relation to existing other scenes. Developers may want to define custom animation settings when the style of a space is changed from one to another.

[3149] The procedure of finding a fitting name for this new concept was a lengthy and challenging one. Instead of calling it a space, next to many other ideas, the following names were more thoroughly considered: Stage — Stage was the internal name used to develop the feature prior to considering API naming. The idea behind this is that ordinary, planar scenes appear on a stage. Due to conflicts with the user-facing Stage Manager feature on phoneOS and computerOS, and the possibility that Stage Control related APIs may be required, and due the hierarchy the word stage implies (scenes semantically happen all on stage, but the API implies rather a peer relationship), the name stage was ultimately discarded.

ImmersiveScene — the idea of including "immersive" in the API name may be appealing for naming APIs specifically for this platform. However, this would rule out bringing the API to other platforms that do not provide any immersion capabilities. Additionally, the suffix "Scene" would be unfortunate, given that existing scene types do not use it either. Panorama — Panorama was considered as a name since it describes the feature quite well, comparing it with the definition of it ("immersive experiences", "360 degree"). Another benefit of panorama is that it is not in use in UI frameworks yet for API naming, and would allow us to expand the API in future further without risking conflicts with existing API naming. However, in order to be closer aligned with marketing and product naming decisions, the name panorama was discarded in favor of space. MixedReality, VirtualReality and VirtualRealityPortal were considered as names for the three styles of immersive spaces. However, to better align with marketing terminology, we renamed these to the currently suggested terms instead.

[3150] Similarly to .formsheet or .popover , the described functionality of a space to display unbounded contents and entering Exclusive Mode could alternatively be achieved via a .immersiveSpace View modifier:

[3151] ```swift

[3152] var body: some View {

[3153] @State private var showingImmersiveSpace: Bool = false

[3154] Button("Show Immersive Space") {

[3155] showingImmersiveSpace.toggle()

[3156] }

[3157] .immersiveSpace(isPresented: \$showingImmersiveSpace, presentation:
 .backdrop) {

[3158] View {

[3159] // ...

[3160] }

[3161] }

[3162] .onAppear { ... }

[3163] .onDisappear { ... }

[3164] }

[3165] ``

[3166]

[3167] One advantage of this could be code locality: The space is defined right where it is presented. However, there are several drawbacks with this direction, which is why we decided against it: The same space and space content may need to be presented from numerous places throughout the application logic, so having a more high-level primitive seems more developer-friendly. Embedding the conceptionally independent view hierarchy of a space locally in an existing view hierarchy seems odd and against separation of responsibilities. Dismissing the view hierarchy from which the space was originally opened is not easily doable. Features such as state restoration, scene phases, and more, are not easily applicable. It became clear that we want ImmersiveSpace to be a first class, prominent, high-level primitive that acts as a building block for the application's user interface just as much as WindowGroup, Window, and other Scene types.

[3168] Another alternative is to create a space by creating a Window instance first, and then apply an immersion style on it, which would give the window the unique space characteristics. However, since we believe that ImmersiveSpace and Window may further diverge in future, we consider it is worth having a dedicated Scene type ImmersiveSpace instead. Note that this means however, that any window modifier can be applied to a space and the space related modifiers can be applied to any other window type. For the combinations that would not make sense, these would simply be no-ops.

[3169] Instead of openImmersiveSpace(...) async -> OpenImmersiveSpace.Result , we also considered openImmersiveSpace() async throws . This however would require the developer to always use try for this callable, but callables have usually been designed as succinct as possible. The benefit of instead indicating success or failure via a return value is that the method can be marked with @discardableResult , so that clients that have no interest

in the result can benefit from a more minimalist code for opening a space. Also note that the set of possible errors in the context of opening a space is rather limited (as opposed to, say, file I/O errors, as in the case for the `openDocument()` async throws callable).

[3170] Instead of `openImmersiveSpace(...)` async -> `OpenImmersiveSpaceAction.Result` , we also considered `openImmersiveSpace()` async -> `Swift.Result<Void, OpenImmersiveSpaceError>` . However, `Swift.Result` requires a value type for the success case, which in the case of `openImmersiveSpace` would be `Void`. We believe that a custom `OpenImmersiveSpaceAction.Result` with three custom tailored cases more straight forward to use for developers (e.g., `.opened` for the success case), than a `Swift.Result` , with which developers would have to compare against `.success(Void)`. The system's default behavior is to hand over a bounded, planar `UIWindowScene` to the application at launch time, with the `UIWindowSceneSessionRoleApplication` scene session role applied. Developers however will want to configure their application to launch with a single immersive space scene, in order to jump directly into an immersive mixed reality or virtual reality experience. Therefore, a static signal is needed for the system-shell to know what kind of scene to initially create. Static because the scene session role may also influence other aspects of the system experience such as scene placeholders, which need to be displayed prior to launch. By introducing a new scene session role for `ImmersiveSpace` scenes, the developer can provide such signal by referring to this role via the new `UIPreferredDefaultSceneSessionRole` key (see `rdar://89832690`) in their application's `Info.plist` . The system-shell will then read the defined value prior to launch, look up the correct scene configuration in the scene manifest, and instantiate a correctly configured scene at launch time. The application is then handed over a scene of the expected type and can embed its user interface. This proposal focuses on the introduction of the scene session role, and its usage for configuring the application's `Info.plist` . Requesting an immersive space scene using the new scene session role at runtime is not covered by this proposal.

[3171] Introduce a new `UISceneSessionRole` called `UISceneSessionRoleImmersiveSpaceApplication` . In Swift, it is exposed as `UISceneSession.Role.immersiveSpaceApplication` . Developers would use the new scene session role in their application's `Info.plist` as a `_value_` for the `UIPreferredDefaultSceneSessionRole` key. Developers may also use the scene session role

as a `_key_` in the dictionary defining the scene configurations their application's `Info.plist`.
Add the following to `<<UIKit/UIScene.h>>`

[3172] ````objective-c`

[3173] `/// A session role for Immersive Space scenes.`

[3174] `UIKIT_EXTERN UISceneSessionRole const`

`UISceneSessionRoleImmersiveSpaceApplication API_AVAILABLE(mixed reality operating system (1.0)) API_UNAVAILABLE(phoneOS, computerOS, mediaOS, wearableOS);`

[3175] `````

[3176] ````swift`

[3177] `extension UISceneSession.Role {`

[3178] `@available(phoneOS, unavailable)`

[3179] `@available(computerOS, unavailable)`

[3180] `@available(mediaOS, unavailable)`

[3181] `@available(wearableOS, unavailable)`

[3182] `@available(mixed reality operating system 1.0, *)`

[3183] `static let immersiveSpaceApplication: UISceneSession.Role`

[3184] `}`

[3185] `````

[3186] This change is purely additive, to enable new behaviors. Existing applications running on will need to ensure that anywhere they check a session role, that they properly handle this new session role in addition to the existing ones. Alternatively to naming it `UISceneSessionRoleImmersiveSpaceApplication`, we could also call it `UIWindowSceneSessionRoleImmersiveSpaceApplication` or `UIImmersiveSpaceSceneSessionRoleImmersiveSpaceApplication`. However, since we likely will not introduce any API for a dedicated immersive space scene in UIKit in Borealis,

the decision whether this class should be a `UIScene` or a `UIWindowScene` subclass has not been made yet, and choosing any of

`UIWindowSceneSessionRoleImmersiveSpaceApplication` or

`UIImmersiveSpaceSceneSessionRoleImmersiveSpaceApplication` now would imply some constraints for the future API design of a class for such scene type. Calling it

`UISceneSessionRoleImmersiveSpaceApplication` for now would ensure that even `_if_` we ever introduce a dedicated class for this new kind of scene, the name would still be valid (as in, the new class would definitely directly or indirectly inherit from `UIScene`).

[3187] Instead of introducing a new scene session role for allowing launching into an immersive experience via a new scene session role, we could also provide no new API around this feature at all, and instead always launch applications with an ordinary `UIWindowScene`. Once launched, developers could then offer a user affordance at runtime to open an immersive space scene. Developers could even destroy the scene session role backing the `UIWindowScene` as soon as they get a chance to together with requesting an immersive space scene, which would come very close to what we are trying to offer, at least from a user experience point of view. However, we considered this as a very unfortunate work-around which we do not want to establish in our developer's minds. Also, `_always_` handing over a regular `UIWindowScene` to the applications could cause unclear situations, if a two-dimensional framework based application uses `ImmersiveSpace` scene as its only scene.

[3188] The system-shell needs `_some_` static signal prior to app launch in order to enable the "launching into an immersive experience" use case. An alternative to providing a new scene session role could be to offer instead a different key-value pair, e.g.,

`UIApplicationLaunchIntoImmersiveExperience`, or something similar. However, in order to pave the way for a future UIKit Immersive Space scene API, we considered it as the right thing to do instead to start off with a dedicated scene session role which hopefully will not have to be deprecated once new immersive space related UIKit APIs will be

introduced. **ImmersiveSpaceDisplacement** On , developers create immersive experiences using a new type of scene that is called `ImmersiveSpace` (see [ImmersiveSpace proposal](<https://github.pie.apple.com/TDG/two-dimensional-framework-Evolution/pull/3>)).

When opening an immersive space scene, by default, the system positions the scene at the

user's feet, or at least somewhere nearby. The developer has no control over the position of the immersive space scene.

[3189] A user may be currently in either a single-user experience, or participate in a shared group experience, using the Group Activities framework. In order to ensure spatial truth during a shared group experience, the system may move the immersive space, and therefore, the general user experience, slightly away from the user's own current location and also rotate it. In this context, we refer with "ego-centric" to the user's own current location in world space, and with "experience-centric" to the center of the (possibly shared) user experience. Accordingly, the two do not have to always match.

[3190] Knowing the difference between the two locations may however be necessary for developers to position content within the immersive space relatively to the user or the experience. Therefore, we propose API to read the pose (a translation and a rotation) that the system has currently applied on the immersive space.

[3191] Introduce a new ``immersiveSpaceDisplacement(in:)`` API on ``GeometryProxy3D``, which is introduced with [this proposal](<https://github.pie.apple.com/TDG/two-dimensional-framework-Evolution/pull/14>).

[3192] Developers use this method to obtain the pose applied by the system to move and rotate the immersive space scene to a more appropriate location. The pose is expressed with ``Pose3D``, which is a structure defined in the Spatial library. The method takes a coordinate space as an argument in order to specify what coordinate space the returned pose should be expressed in.

[3193] A use case of this might look as follows, which displays the content of a game in a immersive space scene, together with a HUD view. The game content appears at the center of the shared group activity ("experience-centric"), whereas the HUD will appear at the ego-centric position by taking the displacement into account.

[3194] ````swift`

[3195] `struct GameApp: App {`

[3196] `var body: some Scene {`

```
[3197]         ImmersiveSpace {
[3198]             GeometryReader3D { proxy in
[3199]                 let displacement =
[3200]                     proxy.immersiveSpaceDisplacement(in: .global)
[3201]                 let inverseDisplacement = displacement.inverse
[3202]             Container3D {
[3203]                 GameContentView()
[3204]                 GameHUDView()
[3205]                 .offset(inverseDisplacement.position)
[3206]                 .rotation3DEffect(inverseDisplacement.rotation)
[3207]             }
[3208]         }
[3209]     }
[3210] }
[3211] }
[3212] ```
[3213]
[3214] ```swift
[3215] extension GeometryProxy3D {
[3216]     /// Returns the pose applied by the system to displace the origin of the
[3217]     /// immersive space. By default, the immersive space is centered near the
```

```

[3218]    /// user's feet, or somewhere close to it. In this case, `.identity` is
[3219]    /// returned. In shared experiences, the center of the experience may be
[3220]    /// moved by the system to a more appropriate location.
[3221]    /// If there is no immersive space currently opened, this returns
[3222]    /// `.identity`.
[3223]    ///
[3224]    /// - Parameters:
[3225]    ///   - space: A coordinate space that the returned pose is relative to
[3226]    @available(phoneOS, unavailable)
[3227]    @available(computerOS, unavailable)
[3228]    @available(mediaOS, unavailable)
[3229]    @available(wearableOS, unavailable)
[3230]    @available(mixed reality operating system 1.0, *)
[3231]    func immersiveSpaceDisplacement(in space: some CoordinateSpaceProtocol)
[3232]        -> Pose3D
[3233]    }
[3234]    ```

```

[3235] Future enhancements may provide access to the poses of other participants of a shared group experience.

[3236] Instead of `experienceDisplacement(in:)`, we also considered other names, but none of them seemed like a perfect fit. In order for this API to stay `ImmersiveSpace` independent, we considered "experience displacement" rather than "immersive space displacement", but the idea was discarded, since "experience" is not something we use

anywhere else yet. Adding a `system` prefix, was but decided against it, given that it is in fact the system that controls the center of the experience,.

[3237] This would no longer require developers to invert the pose before applying it on UI that is supposed to be positioned "ego-centric'ly". However, the naming of such API was even more challenging, since "user pose" could suggest that it updates when the user moves (or changes their head pose), even though it does not – the returned pose really only changes when the Space is displaced, but the user can still move freely to a certain degree with no changes. A clear, none-misleading name for such API could not be found, which is why we stick to the current approach.

[3238] Rather than a `Pose3D`, which consists of a position (`Position3D`) and a rotation (`Rotation3D`), we also considered returning a `AffineTransform3D` instead. However, given that an instance of `AffineTransform3D` also embodies a scale, which is not anything that the system would apply on the overall user experience, we figured that a pose is enough to express the data.

[3239] An alternative to adding a method on `GeometryReader3D` could be to provide the pose applied by the system via the two-dimensional framework environment instead. However, there would be no way to specify the coordinate space relatively to which the pose should be returned.

[3240] Alternatively, we considered introducing a `ImmersiveSpaceReader`, which similarly to `GeometryReader` and `GeometryReader3D`, would provide a proxy object offering a single method similar to the one suggested by this proposal. However, we figured that there is no strong argument introduce such new concept just for the sake of this API, and due to the geometric nature of the new API, the `GeometryReader3D` was considered a good fit.

[3241] The scene manifest, which is an optional part of the application's `Info.plist`, provides a static description of some of the application's aspects that can be indexed and read by the system prior to launching the application. Part of the scene manifest are the scene configurations, which already allow developers to define a name for the configuration, a custom scene class and a custom scene delegate class. The system takes these settings into account during the launch path for setting up the first scene.

[3242] With the introduction of `'UIApplicationPreferredDefaultSceneSessionRole'` key (see `rdar://83741369`), developers can choose which scene session role the system-shell should consider when creating the first scene of the application. If a developer chooses the immersive space scene session role (see `rdar://107833628`), a immersive space scene will be created, which connects to the application at launch. Once an immersive space scene is created by the system-shell, certain aspects of it cannot be changed anymore, including the style of it. Therefore, there is a need for having the necessary configuration readily readable prior to the creation of the immersive space scene.

[3243] We propose to provide a new key, `'UISceneInitialImmersionStyle'`, and new values, `'UIImmersionStyleMixed'`, `'UIImmersionStyleFull'` and `'UIImmersionStyleProgressive'`, that developers can optionally use as part of a scene configuration to configure an immersive space scene statically prior to its creation.

[3244] For an application that wants to launch right into a Virtual Reality user experience, the immersive space scene related `'Info.plist'` declaration may look as follows:

[3245]

[3246] `'''`

[3247] `...`

[3248] `<key>UIApplicationPreferredDefaultSceneSessionRole</key>`

[3249] `<string>UISceneSessionRoleImmersiveSpaceApplication</string>`

[3250] `<key>UISceneConfigurations</key>`

[3251] `<dict>`

[3252] `<key>UISceneSessionRoleImmersiveSpaceApplication</key>`

[3253] `<array>`

[3254] `<dict>`

[3255] `<key>UISceneConfigurationName</key>`

[3256] <string>Immersive Space Configuration</string>

[3257] <key>UISceneInitialImmersionStyle</key>

[3258] <string>UIImmersionStyleFull</string>

[3259] </dict>

[3260] </array>

[3261] ...

[3262] </dict>

[3263] ...

[3264] ``

[3265] Note that in Borealis, two-dimensional framework's `ImmersiveSpace` scene does not allow developers to define a custom scene delegate class nor a custom scene sub-class. Also, Borealis comes with no support for storyboards. Hence, the existing `Info.plist` scene manifest keys `UISceneClassName`, `UISceneStoryboardFile` and `UISceneDelegateClassName` are being ignored by the system for immersive space scenes.

[3266] Also note that this proposal covers the new keys and possible values for the scene manifest section of an application's `Info.plist`. There are no additional symbols for the runtime API necessary.

[3267] Add new optional `Info.plist` keys and a set of allowed, valid values for specifying an immersive space scene configuration.

[3268] | Key | Type | Values

|

[3269] | UISceneInitialImmersionStyle | String | One of {UIImmersionStyleMixed, UIImmersionStyleFull, UIImmersionStyleProgressive} |

[3270] The configuration will be evaluated by the system, but no guarantees can be made about them being honored. If a value cannot be honored, the system will choose a meaningful fallback configuration, and log an error.

[3271] None, since this is a new feature.

[3272] Instead of `'UISceneInitialImmersionStyle'`, a shorter name, `'UIImmersiveSpaceSceneImmersionStyle'`, was also considered. But that name could be misleading, since it suggests that the referenced style is the only style the immersive space scene supports. However, the developer defines the `_initial_` immersive space style with this key, whereas in the developer's two-dimensional framework-based sources, the immersive space may define additionally supported styles, that the immersive space can switch to at runtime.

[3273] Instead of `'UIImmersiveSpaceSceneImmersionStyle'`, `'UISceneImmersiveSpaceStyle'` was also considered, in order to have shared prefix for all keys in a scene manifest's scene configuration. However, a dedicated `'UIImmersiveSpaceScene'` prefix for the key acts as a good indicator that this is a key applicable on an immersive space kind of scene only, and would also align the key better with the naming of the key's possible values.

[3274] We could do nothing, but this would mean that clients could only launch with an immersive space scene with a system-defined default configuration applied, which developers would have no control over. Given that this is a new platform, we do not want to make assumptions about what such default configuration would ideally look like to cover most of the use cases.

[3275] We could allow more key-value based configuration, including hands visibility and video pass-through dimming, but these are aspects of a stage that can also be easily controlled by the developer after the immersive space scene's creation, that is, at runtime. We prefer to start with a minimal set of API and add only what is necessary based on feedback we gather post launch of the platform.

[3276] Instead of a scene configuration, we could also offer other `'Info.plist'` keys and values, (e.g., at the root level of the property list, to define a configuration for the initial immerse space scene). However, given that the concept of scene configurations already

exists, and given the semantic fit for it, and given that there are plans to possibly offer an `ImmersiveSpace` scene API in UIKit post Borealis, using the existing scene configuration pattern for this ask seemed appropriate.

[3277] On mixed reality operating system, developers create immersive experiences using a new type of scene that is called ImmersiveSpace (see ImmersiveSpace proposal). ImmersiveSpace scenes can coexist with regular windowed scenes but there is no way to get the coordinates of any regular view in the coordinate space of the ImmersiveSpace scene. Providing such an API enables an application to move 3D objects in between planar Window scenes and the global ImmersiveSpace and create richer interaction in between the UI of the app and the virtual world.

[3278] Introduce a new NamedCoordinateSpace named `.immersiveSpace`. Developers use this property to obtain the CoordinateSpace corresponding to the currently opened immersive space. `.immersiveSpace` is above `.global` in the hierarchy of coordinate spaces. If there is no currently opened ImmersiveSpace scene in the application, using `.immersiveSpace` for a transform will give the same result as using `.global`

[3279] A possible use case of this API is the following: A developer would like to position an entity part of a view hierarchy in an ImmersiveSpace scene relatively to content that is presented in a WindowGroup scene, or any other window providing scene. For example, a car model that is added to immersive space, in order to be rendered without the limitations of the clipping boundaries of a window, could be positioned close to an inspector window, which allows a user to learn about selected details of the car. A developer would use the new API to obtain the transform of the inspector user interface in the coordinate space of the ImmersiveSpace scene. The transform can then be applied using the right geometry and layout modifiers, to position the car next to the inspector.

[3280] `@available(phoneOS, unavailable)`

[3281] `@available(computerOS, unavailable)`

[3282] `@available(mediaOS, unavailable)`

[3283] `@available(wOS, unavailable)`

[3284] `@available(mixed reality operating system 1.0, *)`

```

[3285]   extension CoordinateSpaceProtocol where Self == NamedCoordinateSpace {
[3286]     /// The named coordinate space that represents the currently opened
[3287]     /// ``ImmersiveSpace`` scene.
[3288]     /// If no immersive space is currently opened, using this CoordinateSpace
[3289]     /// will return the same result as using the ``.global`` coordinate space.
[3290]     public static var immersiveSpace: NamedCoordinateSpace
[3291]   }

```

[3292] We considered the use of a property of GeometryProxy3D instead but we found the use of a CoordinateSpace felt more natural the platform and two-dimensional framework API. We also considered redefining .global but the idea was rejected because the implications for the behaviour of .global on other plaforms like macOS were problematic and we didn't want to create a behaviour exception for mixed reality operating system.

[3293] FIG. 8 illustrates an electronic system 800 with which one or more implementations of the subject technology may be implemented. The electronic system 800 can be, and/or is a part of, the electronic device 105, the handheld electronic device 104, the electronic device 110, the electronic device 115, and/or the server 120 as shown in FIG. 1. The electronic system 800 may include various types of computer readable media and interfaces for various other types of computer readable media. The electronic system 800 includes a bus 808, one or more processing unit(s) 812, a system memory 804 (and/or buffer), a ROM 810, a permanent storage device 802, an input device interface 814, an output device interface 806, and one or more network interfaces 816, or subsets and variations thereof.

[3294] The bus 808 collectively represents all system, peripheral, and chipset buses that communicatively connect the numerous internal devices of the electronic system 1000. In one or more implementations, the bus 808 communicatively connects the one or more processing unit(s) 812 with the ROM 810, the system memory 804, and the permanent storage device 802. From these various memory units, the one or more processing unit(s) 812 retrieves instructions to execute and data to process in order to execute the processes of the subject

disclosure. The one or more processing unit(s) 812 can be a single processor or a multi-core processor in different implementations.

[3295] The ROM 810 stores static data and instructions that are needed by the one or more processing unit(s) 812 and other modules of the electronic system 800. The permanent storage device 802, on the other hand, may be a read-and-write memory device. The permanent storage device 802 may be a non-volatile memory unit that stores instructions and data even when the electronic system 800 is off. In one or more implementations, a mass-storage device (such as a magnetic or optical disk and its corresponding disk drive) may be used as the permanent storage device 802.

[3296] In one or more implementations, a removable storage device (such as a floppy disk, flash drive, and its corresponding disk drive) may be used as the permanent storage device 802. Like the permanent storage device 802, the system memory 804 may be a read-and-write memory device. However, unlike the permanent storage device 802, the system memory 804 may be a volatile read-and-write memory, such as random-access memory. The system memory 804 may store any of the instructions and data that one or more processing unit(s) 812 may need at runtime. In one or more implementations, the processes of the subject disclosure are stored in the system memory 804, the permanent storage device 802, and/or the ROM 810 (which are each implemented as a non-transitory computer-readable medium). From these various memory units, the one or more processing unit(s) 812 retrieves instructions to execute and data to process in order to execute the processes of one or more implementations.

[3297] The bus 808 also connects to the input and output device interfaces 814 and 806. The input device interface 814 enables a user to communicate information and select commands to the electronic system 1000. Input devices that may be used with the input device interface 814 may include, for example, alphanumeric keyboards and pointing devices (also called “cursor control devices”). The output device interface 806 may enable, for example, the display of images generated by electronic system 1000. Output devices that may be used with the output device interface 806 may include, for example, printers and display devices, such as a liquid crystal display (LCD), a light emitting diode (LED) display, an organic light emitting diode (OLED) display, a flexible display, a flat panel display, a solid state display, a projector, or any other device for outputting information. One or more implementations may include devices that function as both input and output devices, such as

a touchscreen. In these implementations, feedback provided to the user can be any form of sensory feedback, such as visual feedback, auditory feedback, or tactile feedback; and input from the user can be received in any form, including acoustic, speech, or tactile input.

[3298] Finally, as shown in FIG. 8, the bus 808 also couples the electronic system 800 to one or more networks and/or to one or more network nodes, such as the electronic device 110 shown in FIG. 1, through the one or more network interface(s) 816. In this manner, the electronic system 800 can be a part of a network of computers (such as a LAN, a wide area network (“WAN”), or an Intranet, or a network of networks, such as the Internet. Any or all components of the electronic system 800 can be used in conjunction with the subject disclosure.

[3299] These functions described above can be implemented in computer software, firmware or hardware. The techniques can be implemented using one or more computer program products. Programmable processors and computers can be included in or packaged as mobile devices. The processes and logic flows can be performed by one or more programmable processors and by one or more programmable logic circuitry. General and special purpose computing devices and storage devices can be interconnected through communication networks.

[3300] Some implementations include electronic components, such as microprocessors, storage and memory that store computer program instructions in a machine-readable or computer-readable medium (also referred to as computer-readable storage media, machine-readable media, or machine-readable storage media). Some examples of such computer-readable media include RAM, ROM, read-only compact discs (CD-ROM), recordable compact discs (CD-R), rewritable compact discs (CD-RW), read-only digital versatile discs (e.g., DVD-ROM, dual-layer DVD-ROM), a variety of recordable/rewritable DVDs (e.g., DVD-RAM, DVD-RW, DVD+RW, etc.), flash memory (e.g., SD cards, mini-SD cards, micro-SD cards, etc.), magnetic and/or solid state hard drives, read-only and recordable Blu-Ray® discs, ultra density optical discs, any other optical or magnetic media, and floppy disks. The computer-readable media can store a computer program that is executable by at least one processing unit and includes sets of instructions for performing various operations. Examples of computer programs or computer code include machine code, such as is produced by a compiler, and files including higher-level code that are executed by a computer, an electronic component, or a microprocessor using an interpreter.

[3301] While the above discussion primarily refers to microprocessor or multi-core processors that execute software, some implementations are performed by one or more integrated circuits, such as application specific integrated circuits (ASICs) or field programmable gate arrays (FPGAs). In some implementations, such integrated circuits execute instructions that are stored on the circuit itself.

[3302] As used in this specification and any claims of this application, the terms “computer”, “server”, “processor”, and “memory” all refer to electronic or other technological devices. These terms exclude people or groups of people. For the purposes of the specification, the terms display or displaying means displaying on an electronic device. As used in this specification and any claims of this application, the terms “computer readable medium” and “computer readable media” are entirely restricted to tangible, physical objects that store information in a form that is readable by a computer. These terms exclude any wireless signals, wired download signals, and any other ephemeral signals.

[3303] To provide for interaction with a user, implementations of the subject matter described in this specification can be implemented on a computer having a display device, e.g., a CRT (cathode ray tube) or LCD (liquid crystal display) monitor, for displaying information to the user and a keyboard and a pointing device, e.g., a mouse or a trackball, by which the user can provide input to the computer. Other kinds of devices can be used to provide for interaction with a user as well; e.g., feedback provided to the user can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile feedback; and input from the user can be received in any form, including acoustic, speech, or tactile input. In addition, a computer can interact with a user by sending documents to and receiving documents from a device that is used by the user; e.g., by sending web pages to a web browser on a user’s client device in response to requests received from the web browser.

[3304] Embodiments of the subject matter described in this specification can be implemented in a computing system that includes a back end component, e.g., as a data server, or that includes a middleware component, e.g., an application server, or that includes a front end component, e.g., a client computer having a graphical user interface or a Web browser through which a user can interact with an implementation of the subject matter described in this specification, or any combination of one or more such back end, middleware, or front end components. The components of the system can be interconnected by any form or medium of digital data communication, e.g., a communication network.

Examples of communication networks include a local area network (“LAN”) and a wide area network (“WAN”), an inter-network (e.g., the Internet), and peer-to-peer networks (e.g., ad hoc peer-to-peer networks).

[3305] The computing system can include clients and servers. A client and server are generally remote from each other and may interact through a communication network. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other. In some examples, a server transmits data (e.g., an HTML page) to a client device (e.g., for purposes of displaying data to and receiving user input from a user interacting with the client device). Data generated at the client device (e.g., a result of the user interaction) can be received from the client device at the server.

[3306] Implementations within the scope of the present disclosure can be partially or entirely realized using a tangible computer-readable storage medium (or multiple tangible computer-readable storage media of one or more types) encoding one or more instructions. The tangible computer-readable storage medium also can be non-transitory in nature.

[3307] The computer-readable storage medium can be any storage medium that can be read, written, or otherwise accessed by a general purpose or special purpose computing device, including any processing electronics and/or processing circuitry capable of executing instructions. For example, without limitation, the computer-readable medium can include any volatile semiconductor memory, such as RAM, DRAM, SRAM, T-RAM, Z-RAM, and TTRAM. The computer-readable medium also can include any non-volatile semiconductor memory, such as ROM, PROM, EPROM, EEPROM, NVRAM, flash, nvSRAM, FeRAM, FeTRAM, MRAM, PRAM, CBRAM, SONOS, RRAM, NRAM, racetrack memory, FJG, and Millipede memory.

[3308] Further, the computer-readable storage medium can include any non-semiconductor memory, such as optical disk storage, magnetic disk storage, magnetic tape, other magnetic storage devices, or any other medium capable of storing one or more instructions. In one or more implementations, the tangible computer-readable storage medium can be directly coupled to a computing device, while in other implementations, the tangible computer-readable storage medium can be indirectly coupled to a computing device, e.g., via

one or more wired connections, one or more wireless connections, or any combination thereof.

[3309] Instructions can be directly executable or can be used to develop executable instructions. For example, instructions can be realized as executable or non-executable machine code or as instructions in a high-level language that can be compiled to produce executable or non-executable machine code. Further, instructions also can be realized as or can include data. Computer-executable instructions also can be organized in any format, including routines, subroutines, programs, data structures, objects, modules, applications, applets, functions, etc. As recognized by those of skill in the art, details including, but not limited to, the number, structure, sequence, and organization of instructions can vary significantly without varying the underlying logic, function, processing, and output.

[3310] While the above discussion primarily refers to microprocessor or multi-core processors that execute software, one or more implementations are performed by one or more integrated circuits, such as ASICs or FPGAs. In one or more implementations, such integrated circuits execute instructions that are stored on the circuit itself.

[3311] Those of skill in the art would appreciate that the various illustrative blocks, modules, elements, components, methods, and algorithms described herein may be implemented as electronic hardware, computer software, or combinations of both. To illustrate this interchangeability of hardware and software, various illustrative blocks, modules, elements, components, methods, and algorithms have been described above generally in terms of their functionality. Whether such functionality is implemented as hardware or software depends upon the particular application and design constraints imposed on the overall system. Skilled artisans may implement the described functionality in varying ways for each particular application. Various components and blocks may be arranged differently (e.g., arranged in a different order, or partitioned in a different way) all without departing from the scope of the subject technology.

[3312] It is understood that any specific order or hierarchy of blocks in the processes disclosed is an illustration of example approaches. Based upon design preferences, it is understood that the specific order or hierarchy of blocks in the processes may be rearranged, or that all illustrated blocks be performed. Any of the blocks may be performed simultaneously. In one or more implementations, multitasking and parallel processing may be

advantageous. Moreover, the separation of various system components in the implementations described above should not be understood as requiring such separation in all implementations, and it should be understood that the described program components and systems can generally be integrated together in a single software product or packaged into multiple software products.

[3313] As used in this specification and any claims of this application, the terms “base station”, “receiver”, “computer”, “server”, “processor”, and “memory” all refer to electronic or other technological devices. These terms exclude people or groups of people. For the purposes of the specification, the terms “display” or “displaying” means displaying on an electronic device.

[3314] As used herein, the phrase “at least one of” preceding a series of items, with the term “and” or “or” to separate any of the items, modifies the list as a whole, rather than each member of the list (i.e., each item). The phrase “at least one of” does not require selection of at least one of each item listed; rather, the phrase allows a meaning that includes at least one of any one of the items, and/or at least one of any combination of the items, and/or at least one of each of the items. By way of example, the phrases “at least one of A, B, and C” or “at least one of A, B, or C” each refer to only A, only B, or only C; any combination of A, B, and C; and/or at least one of each of A, B, and C.

[3315] The predicate words “configured to”, “operable to”, and “programmed to” do not imply any particular tangible or intangible modification of a subject, but, rather, are intended to be used interchangeably. In one or more implementations, a processor configured to monitor and control an operation or a component may also mean the processor being programmed to monitor and control the operation or the processor being operable to monitor and control the operation. Likewise, a processor configured to execute code can be construed as a processor programmed to execute code or operable to execute code.

[3316] Phrases such as an aspect, the aspect, another aspect, some aspects, one or more aspects, an implementation, the implementation, another implementation, some implementations, one or more implementations, an embodiment, the embodiment, another embodiment, some implementations, one or more implementations, a configuration, the configuration, another configuration, some configurations, one or more configurations, the subject technology, the disclosure, the present disclosure, other variations thereof and alike

are for convenience and do not imply that a disclosure relating to such phrase(s) is essential to the subject technology or that such disclosure applies to all configurations of the subject technology. A disclosure relating to such phrase(s) may apply to all configurations, or one or more configurations. A disclosure relating to such phrase(s) may provide one or more examples. A phrase such as an aspect or some aspects may refer to one or more aspects and vice versa, and this applies similarly to other foregoing phrases.

[3317] The word “exemplary” is used herein to mean “serving as an example, instance, or illustration”. Any embodiment described herein as “exemplary” or as an “example” is not necessarily to be construed as preferred or advantageous over other implementations. Furthermore, to the extent that the term “include”, “have”, or the like is used in the description or the claims, such term is intended to be inclusive in a manner similar to the term “comprise” as “comprise” is interpreted when employed as a transitional word in a claim.

[3318] All structural and functional equivalents to the elements of the various aspects described throughout this disclosure that are known or later come to be known to those of ordinary skill in the art are expressly incorporated herein by reference and are intended to be encompassed by the claims. Moreover, nothing disclosed herein is intended to be dedicated to the public regardless of whether such disclosure is explicitly recited in the claims. No claim element is to be construed under the provisions of 35 U.S.C. § 112(f) unless the element is expressly recited using the phrase “means for” or, in the case of a method claim, the element is recited using the phrase “step for”.

[3319] The previous description is provided to enable any person skilled in the art to practice the various aspects described herein. Various modifications to these aspects will be readily apparent to those skilled in the art, and the generic principles defined herein may be applied to other aspects. Thus, the claims are not intended to be limited to the aspects shown herein, but are to be accorded the full scope consistent with the language claims, wherein reference to an element in the singular is not intended to mean “one and only one” unless specifically so stated, but rather “one or more”. Unless specifically stated otherwise, the term “some” refers to one or more. Pronouns in the masculine (e.g., his) include the feminine and neuter gender (e.g., her and its) and vice versa. Headings and subheadings, if any, are used for convenience only and do not limit the subject disclosure.

[3320] The foregoing description, for purpose of explanation, has been described with reference to specific examples. However, the illustrative discussions above are not intended to be exhaustive or to limit the disclosure to the precise forms disclosed. Many modifications and variations are possible in view of the above teachings. The examples were chosen and described in order to best explain the principles of the techniques and their practical applications. Others skilled in the art are thereby enabled to best utilize the techniques and various examples with various modifications as are suited to the particular use contemplated.

[3321] Although the disclosure and examples have been fully described with reference to the accompanying drawings, it is to be noted that various changes and modifications will become apparent to those skilled in the art. Such changes and modifications are to be understood as being included within the scope of the disclosure and examples as defined by the claims.

[3322] As described above, one aspect of the present technology is the gathering and use of data available from various sources to improve how a device interacts with a user. The present disclosure contemplates that in some instances, this gathered data can include personal information data that uniquely identifies or can be used to contact or locate a specific person. Such personal information data can include demographic data, location-based data, telephone numbers, email addresses, home addresses, or any other identifying information.

[3323] The present disclosure recognizes that the use of such personal information data, in the present technology, can be used to the benefit of users. For example, the personal information data can be used to change how a device interacts with a user. Accordingly, use of such personal information data enables better user interactions. Further, other uses for personal information data that benefit the user are also contemplated by the present disclosure.

[3324] The present disclosure further contemplates that the entities responsible for the collection, analysis, disclosure, transfer, storage, or other use of such personal information data will comply with well-established privacy policies and/or privacy practices. In particular, such entities should implement and consistently use privacy policies and practices that are generally recognized as meeting or exceeding industry or governmental requirements for maintaining personal information data private and secure. For example, personal

information from users should be collected for legitimate and reasonable uses of the entity and not shared or sold outside of those legitimate uses. Further, such collection should occur only after receiving the informed consent of the users. Additionally, such entities would take any needed steps for safeguarding and securing access to such personal information data and ensuring that others with access to the personal information data adhere to their privacy policies and procedures. Further, such entities can subject themselves to evaluation by third parties to certify their adherence to widely accepted privacy policies and practices.

[3325] Despite the foregoing, the present disclosure also contemplates examples in which users selectively block the use of, or access to, personal information data. That is, the present disclosure contemplates that hardware and/or software elements can be provided to prevent or block access to such personal information data. For example, in the case of image capture, the present technology can be configured to allow users to select to "opt in" or "opt out" of participation in the collection of personal information data during registration for services.

[3326] Therefore, although the present disclosure broadly covers use of personal information data to implement one or more various disclosed examples, the present disclosure also contemplates that the various examples can also be implemented without the need for accessing such personal information data. That is, the various examples of the present technology are not rendered inoperable due to the lack of all or a portion of such personal information data. For example, content can be displayed to users by inferring location based on non-personal information data or a bare minimum amount of personal information, such as the content being requested by the device associated with a user or other non-personal information.

CLAIMS

What is claimed is:

1. A method, comprising:
receiving, from a first application, a request to render a first object in an environment;
and
in response to receiving the request to render the first object in the environment,
rendering the first object;
after rendering the first object and while the first object is being displayed in the
environment, detecting an indication of a request to change a size of content provided by the
first application; and
in response to detecting the indication of the request to change the size of content
provided by the first application, rendering a representation of the environment by:
in accordance with a determination that the environment is rendered according
to a first mode and that the first object is a first type of object, wherein the environment
corresponds to the first application and a second application different from the first
application while the environment is rendered according to the first mode, maintaining a
physical size of the first object without maintaining an angular size of the first object;
in accordance with a determination that the environment is rendered according
to the first mode and that the first object is a second type of object different from the first
type of object, maintaining the angular size of the first object without maintaining the
physical size of the first object; and
in accordance with a determination that the environment is rendered according
to a second mode, wherein the environment corresponds to the first application without
corresponding to another application different from the first application while the
environment is rendered according to the second mode, maintaining the physical size of the
first object without maintaining the angular size of the first object.
2. The method of claim 1, wherein the second mode is a window mode, and wherein the
first object is within a window while the environment is rendered according to the second
mode.
3. The method of any one of claims 1-2, further comprising:

while a second object is being displayed in the environment, detecting a second indication of a request to change a size of content, wherein the second indication is different from the indication; and

in response to detecting the second indication of the request to change the size of content, rendering a second representation of the environment by:

in accordance with a determination that the environment is rendered according to the second mode and that the second object is a third type of object, maintaining a physical size of the second object without maintaining an angular size of the second object; and

in accordance with a determination that the environment is rendered according to the second mode and that the second object is a fourth type of object different from the first type of object, maintaining the angular size of the first object without maintaining the physical size of the first object, wherein the fourth type of object is different from the third type of object.

4. The method of any one of claims 1-3, while a third object is being displayed concurrently with the first object in the environment and in response to detecting the indication of the request to change the size of content provided by the first application, rendering the representation of the environment by:

while maintaining the angular size of the first object, maintaining a physical size of the third object without maintaining an angular size of the third object; and

while maintaining the physical size of the first object, maintaining the angular size of the third object without maintaining the physical size of the third object.

5. The method of any one of claims 1-4, further comprising:

while a fourth object of a second application is being displayed concurrently with the first object in the environment and in response to detecting an indication of a request to change a size of content, wherein the second application is different from the first application, rendering the representation of the environment by:

while maintaining the angular size of the first object, maintaining a physical size of the fourth object without maintaining an angular size of the fourth object; and

while maintaining the physical size of the first object, maintaining the angular size of the fourth object without maintaining the physical size of the fourth object.

6. The method of claim 5, wherein the content provided by the first application is concurrently displayed in the environment with content provided by the second application.
7. The method of any one of claims 1-6, further comprising:
 - while the environment is rendered according to the first mode, detecting a request to transition to the second mode; and
 - in response to detecting the request to transition to the second mode, causing the environment to operate in the second mode, wherein causing the environment to operate in the second mode includes removing, from the environment, content provided by another application different from the first application.
8. The method of any one of claims 1-7, wherein, while the environment is rendered according to the first mode: content corresponding to the first application is in a first area of the environment and not in a second area of the environment and content corresponding to a second application is in the first area and not in the second area, wherein the second area is different from the first area, and wherein the second application is different from the first application.
9. The method of any one of claims 1-8, wherein, while the environment is rendered according to the first mode and in accordance with a determination that a first set of one or more criteria is satisfied, content provided by the first application is within a three-dimensional bounded window.
10. The method of any one of claims 1-9, wherein, while in the first mode and in accordance with a determination that a second set of one or more criteria is satisfied, content corresponding to the first application is within a two-dimensional bounded window, and wherein the second set of one or more criteria is different from the first set of one or more criteria.
11. The method of claim 10, wherein the first set of one or more criteria includes a criterion that is satisfied when a first window corresponding to the first application includes a style defining that the window is a three-dimensional bounded window, and wherein the second set of one or more criteria includes a criterion that is satisfied when a second window

corresponding to the first application includes a style defining that the second window is a two-dimensional bounded window.

12. The method of any one of claims 8-11, wherein the second area surrounds the first area.

13. The method of any one of claims 1-12, wherein:

maintaining the physical size of the first object without maintaining the angular size of the first object includes:

at a first time while the size of content provided by the first object is changed, maintaining the physical size of the first object and maintaining the angular size of the first object.

14. The method of claim 13, wherein:

maintaining the physical size of the first object without maintaining the angular size of the first object includes:

at a second time, different from the first time, while the size of content provided by the first object is changed, changing the angular size of the first object while maintaining the physical size of the first object.

15. A non-transitory computer-readable medium storing one or more programs configured to be executed by one or more processors of a computer system, the one or more programs including instructions for performing the method of any one of claims 1-14.

16. A computer system, comprising:

one or more processors; and

memory storing one or more programs configured to be executed by the one or more processors, the one or more programs including instructions for performing the method of any one of claims 1-14.

17. A computer system, comprising:

means for performing the method of any one of claims 1-14.

18. A computer program product, comprising one or more programs configured to be executed by one or more processors of a computer system, the one or more programs including instructions for performing the method of any one of claims 1-14.

19. A non-transitory computer-readable storage medium storing one or more programs configured to be executed by one or more processors of a computer system, the one or more programs including instructions for:

receiving, from a first application, a request to render a first object in an environment;
and

in response to receiving the request to render the first object in the environment,
rendering the first object;

after rendering the first object and while the first object is being displayed in the environment, detecting an indication of a request to change a size of content provided by the first application; and

in response to detecting the indication of the request to change the size of content provided by the first application, rendering a representation of the environment by:

in accordance with a determination that the environment is rendered according to a first mode and that the first object is a first type of object, wherein the environment corresponds to the first application and a second application different from the first application while the environment is rendered according to the first mode, maintaining a physical size of the first object without maintaining an angular size of the first object;

in accordance with a determination that the environment is rendered according to the first mode and that the first object is a second type of object different from the first type of object, maintaining the angular size of the first object without maintaining the physical size of the first object; and

in accordance with a determination that the environment is rendered according to a second mode, wherein the environment corresponds to the first application without corresponding to another application different from the first application while the environment is rendered according to the second mode, maintaining the physical size of the first object without maintaining the angular size of the first object.

20. A computer system, comprising:
one or more processors; and

memory storing one or more programs configured to be executed by the one or more processors, the one or more programs including instructions for:

receiving, from a first application, a request to render a first object in an environment; and

in response to receiving the request to render the first object in the environment, rendering the first object;

after rendering the first object and while the first object is being displayed in the environment, detecting an indication of a request to change a size of content provided by the first application; and

in response to detecting the indication of the request to change the size of content provided by the first application, rendering a representation of the environment by:

in accordance with a determination that the environment is rendered according to a first mode and that the first object is a first type of object, wherein the environment corresponds to the first application and a second application different from the first application while the environment is rendered according to the first mode, maintaining a physical size of the first object without maintaining an angular size of the first object;

in accordance with a determination that the environment is rendered according to the first mode and that the first object is a second type of object different from the first type of object, maintaining the angular size of the first object without maintaining the physical size of the first object; and

in accordance with a determination that the environment is rendered according to a second mode, wherein the environment corresponds to the first application without corresponding to another application different from the first application while the environment is rendered according to the second mode, maintaining the physical size of the first object without maintaining the angular size of the first object.

21. A computer system, comprising:

means for, receiving, from a first application, a request to render a first object in an environment; and

means for, in response to receiving the request to render the first object in the environment, rendering the first object;

means for, after rendering the first object and while the first object is being displayed in the environment, detecting an indication of a request to change a size of content provided by the first application; and

in response to detecting the indication of the request to change the size of content provided by the first application, rendering a representation of the environment by:

means for, in accordance with a determination that the environment is rendered according to a first mode and that the first object is a first type of object, wherein the environment corresponds to the first application and a second application different from the first application while the environment is rendered according to the first mode, maintaining a physical size of the first object without maintaining an angular size of the first object;

means for, in accordance with a determination that the environment is rendered according to the first mode and that the first object is a second type of object different from the first type of object, maintaining the angular size of the first object without maintaining the physical size of the first object; and

means for, in accordance with a determination that the environment is rendered according to a second mode, wherein the environment corresponds to the first application without corresponding to another application different from the first application while the environment is rendered according to the second mode, maintaining the physical size of the first object without maintaining the angular size of the first object.

22. A computer program product, comprising one or more programs configured to be executed by one or more processors of a computer system, the one or more programs including instructions for:

receiving, from a first application, a request to render a first object in an environment; and

in response to receiving the request to render the first object in the environment, rendering the first object;

after rendering the first object and while the first object is being displayed in the environment, detecting an indication of a request to change a size of content provided by the first application; and

in response to detecting the indication of the request to change the size of content provided by the first application, rendering a representation of the environment by:

in accordance with a determination that the environment is rendered according to a first mode and that the first object is a first type of object, wherein the environment corresponds to the first application and a second application different from the first application while the environment is rendered according to the first mode, maintaining a physical size of the first object without maintaining an angular size of the first object;

in accordance with a determination that the environment is rendered according to the first mode and that the first object is a second type of object different from the first type of object, maintaining the angular size of the first object without maintaining the physical size of the first object; and

in accordance with a determination that the environment is rendered according to a second mode, wherein the environment corresponds to the first application without corresponding to another application different from the first application while the environment is rendered according to the second mode, maintaining the physical size of the first object without maintaining the angular size of the first object.

23. A method, comprising:

sending, from an application, a request for a pose associated with a user with respect to an origin of an experience for the application; and

after sending the request for the pose associated with the user with respect to the origin of the experience for the application:

in accordance with a determination that a first number of users are in the experience for the application:

receiving a first pose; and

rendering the experience for the application based on the first pose;

and

in accordance with a determination that a second number of users are in the experience for the application:

receiving a second pose different from the first pose; and

rendering the experience for the application based on the second pose.

24. The method of claim 23, wherein:

the first number of users is one;

the first pose corresponds to a center of the experience for the application;

the second number of users is two or more; and
the second pose corresponds to an offset greater than zero from the center of the experience for the application.

25. The method of any one of claims 23-24, further comprising:

after sending the request for the pose associated with the user with respect to the origin of the experience for the application and in accordance with a determination that a third number of users are in the experience for the application, wherein the third number of users is different from the first number of users and the second number of users:

receiving a third pose different from the first pose and the second pose; and
rendering the experience for the application based on the third pose.

26. The method of any one of claims 23-25, wherein the second pose is generated based on a plurality of users being placed around an origin of the experience for the application.

27. The method of any one of claims 23-26, wherein the second pose is generated based on multiple users being placed in a sequence near a second origin of the experience for the application.

28. The method of any one of claims 23-27, wherein the request for the pose is sent in response to detecting intent of a user to launch the application.

29. The method of claim 28, wherein, after detecting intent of the user to launch the application in a respective mode, a process of a computer system executing the application performs one or more pre-launch tasks on behalf of the application, wherein the one or more pre-launch tasks correspond to the application launching in the respective mode, and wherein the one or more pre-launch tasks occur before an input is detected to launch the application in the respective mode.

30. The method of any one of claims 23-29, wherein the request for the pose is sent in conjunction with a request to change from a second experience to the experience, wherein the second experience is for the application, and wherein the second experience is different from the experience.

31. The method of any one of claims 23-30, wherein the first pose includes a first displacement from a third origin of the experience, and wherein the second pose includes a second displacement from the third origin of the experience.
32. The method of any one of claims 23-31, wherein the request for the pose includes an indication of a coordinate space.
33. The method of claim 32, wherein:
in accordance with a determination that the request for the pose includes an indication of a first coordinate space, a respective pose that is received corresponds to the first coordinate space; and
in accordance with a determination that the request for the pose includes an indication of a second coordinate space different from the first coordinate space, the respective pose that is received corresponds to the second coordinate space.
34. The method of any one of claims 23-33, further comprising:
after receiving the first pose, sending a second request for the pose associated with the user; and
after sending the second request for the pose associated with the user, receiving one or more affine transformations instead of a pose, wherein the one or more affine transformations correspond to a previous pose received by the application.
35. A non-transitory computer-readable medium storing one or more programs configured to be executed by one or more processors of a computer system, the one or more programs including instructions for performing the method of any one of claims 23-34.
36. A computer system, comprising:
one or more processors; and
memory storing one or more programs configured to be executed by the one or more processors, the one or more programs including instructions for performing the method of any one of claims 23-34.

37. A computer system, comprising:
means for performing the method of any one of claims 23-34.
38. A computer program product, comprising one or more programs configured to be executed by one or more processors of a computer system, the one or more programs including instructions for performing the method of any one of claims 23-34.
39. A non-transitory computer-readable storage medium storing one or more programs configured to be executed by one or more processors of a computer system, the one or more programs including instructions for:
sending, from an application, a request for a pose associated with a user with respect to an origin of an experience for the application; and
after sending the request for the pose associated with the user with respect to the origin of the experience for the application:
in accordance with a determination that a first number of users are in the experience for the application:
receiving a first pose; and
rendering the experience for the application based on the first pose;
and
in accordance with a determination that a second number of users are in the experience for the application:
receiving a second pose different from the first pose; and
rendering the experience for the application based on the second pose.
40. A computer system, comprising:
one or more processors; and
memory storing one or more programs configured to be executed by the one or more processors, the one or more programs including instructions for:
sending, from an application, a request for a pose associated with a user with respect to an origin of an experience for the application; and
after sending the request for the pose associated with the user with respect to the origin of the experience for the application:
in accordance with a determination that a first number of users are in the experience for the application:

receiving a first pose; and
rendering the experience for the application based on the first pose; and

in accordance with a determination that a second number of users are in the experience for the application:

receiving a second pose different from the first pose; and
rendering the experience for the application based on the second pose.

41. A computer system, comprising:

means for, sending, from an application, a request for a pose associated with a user with respect to an origin of an experience for the application; and

after sending the request for the pose associated with the user with respect to the origin of the experience for the application:

in accordance with a determination that a first number of users are in the experience for the application:

means for receiving a first pose; and
means for rendering the experience for the application based on the first pose; and

in accordance with a determination that a second number of users are in the experience for the application:

means for receiving a second pose different from the first pose; and
means for rendering the experience for the application based on the second pose.

42. A computer program product, comprising one or more programs configured to be executed by one or more processors of a computer system, the one or more programs including instructions for:

sending, from an application, a request for a pose associated with a user with respect to an origin of an experience for the application; and

after sending the request for the pose associated with the user with respect to the origin of the experience for the application:

in accordance with a determination that a first number of users are in the experience for the application:

receiving a first pose; and

rendering the experience for the application based on the first pose;

and

in accordance with a determination that a second number of users are in the experience for the application:

receiving a second pose different from the first pose; and

rendering the experience for the application based on the second pose.








FIG. 1

**FIG. 2**

**FIG. 3**

**FIG. 4**

**FIG. 5**




FIG. 6A





FIG. 6B

**FIG. 7**




FIG. 8

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2024/030462

A. CLASSIFICATION OF SUBJECT MATTER					
INV.	G06F9/451	G06F3/04815	G06F3/04845	G06T19/00	G06T19/20
ADD.	G06F3/01	G06F3/0481	G06F3/04842		
According to International Patent Classification (IPC) or to both national classification and IPC					
B. FIELDS SEARCHED					
Minimum documentation searched (classification system followed by classification symbols) G06F G06V G06T					
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, COMPENDEX, INSPEC, IBM-TDB, WPI Data					
C. DOCUMENTS CONSIDERED TO BE RELEVANT					
Category*	Citation of document, with indication, where appropriate, of the relevant passages				Relevant to claim No.
X	US 2023/094522 A1 (STAUBER MATAN [US] ET AL) 30 March 2023 (2023-03-30)				1-12, 15-22
Y	par.125, 129, 131-133, 137, 245-246, 250, 253, 255; figures 7A-D, 9A, 10C -----				13,14
Y	US 2023/152935 A1 (MCKENZIE CHRISTOPHER D [US] ET AL) 18 May 2023 (2023-05-18)				13,14
A	par.399-400 -----				1-12, 15-22
☐ Further documents are listed in the continuation of Box C. ☒ See patent family annex.					
* Special categories of cited documents : "A" document defining the general state of the art which is not considered to be of particular relevance "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 1 August 2024			Date of mailing of the international search report 30/09/2024		
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 Filimon, Diana		

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2024/030462

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

see additional sheet

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☒ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims;; it is covered by claims Nos.:

1 - 22

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

This International Searching Authority found multiple (groups of) inventions in this international application, as follows:

1. claims: 1-22

Managing size of virtual objects according to their type, while a content rendered in a virtual reality environment including the virtual objects is resized, to ensure appropriate perception of objects.

2. claims: 23-42

Manage user position in computer generated virtual reality experiences according to number of users in the experience.

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2024/030462

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2023094522 A1	30-03-2023	AU 2022348895 A1	11-04-2024
		CN 118159935 A	07-06-2024
		EP 4405797 A1	31-07-2024
		KR 20240065152 A	14-05-2024
		US 2023094522 A1	30-03-2023
		WO 2023049671 A1	30-03-2023

US 2023152935 A1	18-05-2023	AU 2022352933 A1	11-04-2024
		CN 118215903 A	18-06-2024
		EP 4388397 A1	26-06-2024
		KR 20240064014 A	10-05-2024
		US 2023152935 A1	18-05-2023
		WO 2023049670 A1	30-03-2023
