



- (51) International Patent Classification:  
*G06F 3/01* (2006.01)      *G06T 19/20* (2011.01)  
*G06T 19/00* (2011.01)

(21) International Application Number:  
PCT/US2024/030181

(22) International Filing Date:  
20 May 2024 (20.05.2024)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
63/471,209      05 June 2023 (05.06.2023)      US  
18/622,447      29 March 2024 (29.03.2024)      US

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

(72) Inventors: **RICHARDSON, Andrew P.**; c/o Apple Inc., One Apple Park Way, Cupertino, California 95014 (US). **NILES, Christian A.**; c/o Apple Inc., One Apple Park Way, Cupertino, California 95014 (US). **RUSSELL, Collin R.**; c/o Apple Inc., One Apple Park Way, Cupertino, California 95014 (US). **ASHUTOSH, Abhinay**; c/o Apple Inc.,
- One Apple Park Way, Cupertino, California 95014 (US). **EBBOLE, Mark A.**; 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,

(54) Title: TECHNIQUES FOR THREE-DIMENSIONAL ENVIRONMENTS




FIG. 4

(57) Abstract: Some techniques are described herein for integrating a 2D framework with a 3D framework. Such techniques use a concept referred to as a hidden entity to link the two frameworks together. Other techniques are described herein for translating gestures from a first type to a second type in certain situations.

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,  
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 THREE-DIMENSIONAL ENVIRONMENTS

### CROSS-REFERENCE TO RELATED APPLICATIONS

**[0001]** The present application claims priority to U.S. Non-Provisional Patent Application Serial No. 18/622,447, entitled “TECHNIQUES FOR THREE-DIMENSIONAL ENVIRONMENTS” filed March 29, 2024, and to U.S. Provisional Patent Application Serial No. 63/471,209, entitled “TECHNIQUES FOR THREE-DIMENSIONAL ENVIRONMENTS” filed June 5, 2023, which are hereby incorporated by reference in their entireties for all purposes.

### BACKGROUND

**[0002]** Electronic devices are becoming increasingly complex. For example, a single electronic device can utilize multiple frameworks for different situations. Ensuring that such frameworks work together and/or are able to be used together can be difficult. Accordingly, there is a need to improve techniques for facilitating operations.

### SUMMARY

**[0003]** Current techniques for managing three-dimensional (3D) environments are generally ineffective and/or inefficient. For example, some techniques require frameworks for two-dimensional (2D) environments to be kept separate from frameworks for three-dimensional environments. Other techniques avoid this issue altogether and attempt to use a one-size fits all model that can result in issues. This disclosure provides more effective and/or efficient techniques for managing three-dimensional environments using an example of a single computer system. It should be recognized that other types of electronic devices can be used with techniques described herein. For example, a computer system relying on remote rendering of content can use techniques described herein. In addition, techniques optionally complement or replace other techniques for managing three-dimensional environments.

**[0004]** Some techniques are described herein for integrating a 2D framework with a 3D framework. Such techniques use a concept referred to as a hidden entity to link the two frameworks together. Other techniques are described herein for translating gestures from a first type to a second type in certain situations.

**[0005]** In some examples, a method that is performed by a computer system is described. In some examples, the method comprises: receiving, from an application, a request to add a two-dimensional entity at a first location in a three-dimensional environment, wherein the 3D environment includes one or more 3D entities; and in response to receiving the request to add the 2D entity at the first location in the 3D environment: adding a first 3D entity to the first location in the 3D environment; rendering, via a 2D framework, a representation of the 2D entity; and rendering, via a 3D framework, a representation of a second 3D entity by: performing one or more operations on the representation of the 2D entity; and placing the representation of the second 3D entity at the first location.

**[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 an application, a request to add a two-dimensional entity at a first location in a three-dimensional environment, wherein the 3D environment includes one or more 3D entities; and in response to receiving the request to add the 2D entity at the first location in the 3D environment: adding a first 3D entity to the first location in the 3D environment; rendering, via a 2D framework, a representation of the 2D entity; and rendering, via a 3D framework, a representation of a second 3D entity by: performing one or more operations on the representation of the 2D entity; and placing the representation of the second 3D entity at the first location.

**[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 an application, a request to add a two-dimensional entity at a first location in a three-dimensional environment, wherein the 3D environment includes one or more 3D entities; and in response to receiving the request to add the 2D entity at the first location in the 3D environment: adding a first 3D entity to the first location in the 3D environment; rendering, via a 2D framework, a representation of the 2D entity; and rendering, via a 3D framework, a representation of a second 3D entity by: performing one or more operations on the representation of the 2D entity; and placing the representation of the second 3D entity at the first location.

**[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 an application, a request to add a two-dimensional entity at a first location in a three-dimensional environment, wherein the 3D environment includes one or more 3D entities; and in response to receiving the request to add the 2D entity at the first location in the 3D environment: adding a first 3D entity to the first location in the 3D environment; rendering, via a 2D framework, a representation of the 2D entity; and rendering, via a 3D framework, a representation of a second 3D entity by: performing one or more operations on the representation of the 2D entity; and placing the representation of the second 3D entity at the first location.

**[0009]** In some examples, a computer system is comprising means for performing each of the following steps: receiving, from an application, a request to add a two-dimensional entity at a first location in a three-dimensional environment, wherein the 3D environment includes one or more 3D entities; and in response to receiving the request to add the 2D entity at the first location in the 3D environment: adding a first 3D entity to the first location in the 3D environment; rendering, via a 2D framework, a representation of the 2D entity; and rendering, via a 3D framework, a representation of a second 3D entity by: performing one or more operations on the representation of the 2D entity; and placing the representation of the second 3D entity at the first location.

**[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 an application, a request to add a two-dimensional entity at a first location in a three-dimensional environment, wherein the 3D environment includes one or more 3D entities; and in response to receiving the request to add the 2D entity at the first location in the 3D environment: adding a first 3D entity to the first location in the 3D environment; rendering, via a 2D framework, a representation of the 2D entity; and rendering, via a 3D framework, a representation of a second 3D entity by: performing one or more operations on the representation of the 2D entity; and placing the representation of the second 3D entity at the first location.

**[0011]** In some examples, a method that is performed at a computer system in communication with one or more input devices is described. In some examples, the method comprises: detecting, via the one or more input devices, a first input corresponding to a three-dimensional environment; in response to detecting the first input corresponding to the 3D environment: in accordance with a determination that a first set of one or more criteria is satisfied, wherein the first set of one or more criteria includes a criterion that is satisfied when the input is directed to a first entity of a first type in the 3D environment: translating the first input to a second input different from the first input; and sending, to a first application, an indication of the second input; and in accordance with a determination that a second set of one or more criteria is satisfied, wherein the second set of one or more criteria includes a criterion that is satisfied when the input is directed to a second entity of a second type in the 3D environment, sending, to a second application, an indication of the first input, wherein the second type of entity is different from the first type of entity, and wherein the second set of one or more criteria is different from the first set of one or more criteria.

**[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 in communication with one or more input devices is described. In some examples, the one or more programs includes instructions for: detecting, via the one or more input devices, a first input corresponding to a three-dimensional environment; in response to detecting the first input corresponding to the 3D environment: in accordance with a determination that a first set of one or more criteria is satisfied, wherein the first set of one or more criteria includes a criterion that is satisfied when the input is directed to a first entity of a first type in the 3D environment: translating the first input to a second input different from the first input; and sending, to a first application, an indication of the second input; and in accordance with a determination that a second set of one or more criteria is satisfied, wherein the second set of one or more criteria includes a criterion that is satisfied when the input is directed to a second entity of a second type in the 3D environment, sending, to a second application, an indication of the first input, wherein the second type of entity is different from the first type of entity, and wherein the second set of one or more criteria is different from the first set of one or more criteria.

**[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 in

communication with one or more input devices is described. In some examples, the one or more programs includes instructions for: detecting, via the one or more input devices, a first input corresponding to a three-dimensional environment; in response to detecting the first input corresponding to the 3D environment: in accordance with a determination that a first set of one or more criteria is satisfied, wherein the first set of one or more criteria includes a criterion that is satisfied when the input is directed to a first entity of a first type in the 3D environment: translating the first input to a second input different from the first input; and sending, to a first application, an indication of the second input; and in accordance with a determination that a second set of one or more criteria is satisfied, wherein the second set of one or more criteria includes a criterion that is satisfied when the input is directed to a second entity of a second type in the 3D environment, sending, to a second application, an indication of the first input, wherein the second type of entity is different from the first type of entity, and wherein the second set of one or more criteria is different from the first set of one or more criteria.

**[0014]** In some examples, a computer system in communication with one or more input devices is described. In some examples, the computer system in communication with one or more input devices comprises one or more processors and memory storing one or more programs configured to be executed by the one or more processors. In some examples, the one or more programs includes instructions for: detecting, via the one or more input devices, a first input corresponding to a three-dimensional environment; in response to detecting the first input corresponding to the 3D environment: in accordance with a determination that a first set of one or more criteria is satisfied, wherein the first set of one or more criteria includes a criterion that is satisfied when the input is directed to a first entity of a first type in the 3D environment: translating the first input to a second input different from the first input; and sending, to a first application, an indication of the second input; and in accordance with a determination that a second set of one or more criteria is satisfied, wherein the second set of one or more criteria includes a criterion that is satisfied when the input is directed to a second entity of a second type in the 3D environment, sending, to a second application, an indication of the first input, wherein the second type of entity is different from the first type of entity, and wherein the second set of one or more criteria is different from the first set of one or more criteria.

**[0015]** In some examples, a computer system in communication with one or more input devices is described. In some examples, the computer system in communication with one or more input devices comprises means for performing each of the following steps: detecting, via the one or more input devices, a first input corresponding to a three-dimensional environment; in response to detecting the first input corresponding to the 3D environment: in accordance with a determination that a first set of one or more criteria is satisfied, wherein the first set of one or more criteria includes a criterion that is satisfied when the input is directed to a first entity of a first type in the 3D environment: translating the first input to a second input different from the first input; and sending, to a first application, an indication of the second input; and in accordance with a determination that a second set of one or more criteria is satisfied, wherein the second set of one or more criteria includes a criterion that is satisfied when the input is directed to a second entity of a second type in the 3D environment, sending, to a second application, an indication of the first input, wherein the second type of entity is different from the first type of entity, and wherein the second set of one or more criteria is different from the first set of one or more criteria.

**[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 communication with one or more input devices. In some examples, the one or more programs include instructions for: detecting, via the one or more input devices, a first input corresponding to a three-dimensional environment; in response to detecting the first input corresponding to the 3D environment: in accordance with a determination that a first set of one or more criteria is satisfied, wherein the first set of one or more criteria includes a criterion that is satisfied when the input is directed to a first entity of a first type in the 3D environment: translating the first input to a second input different from the first input; and sending, to a first application, an indication of the second input; and in accordance with a determination that a second set of one or more criteria is satisfied, wherein the second set of one or more criteria includes a criterion that is satisfied when the input is directed to a second entity of a second type in the 3D environment, sending, to a second application, an indication of the first input, wherein the second type of entity is different from the first type of entity, and wherein the second set of one or more criteria is different from the first set of one or more criteria.

[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 of a scene graph in accordance with some examples.

[0023] FIG. 5 is a block diagram illustrating a process for processing gestures in a 3D environment according to some examples.

[0024] FIG. 6 is a flow diagram illustrating a method for integrating a 2D framework with a 3D framework in accordance with some examples.

[0025] FIG. 7 is a flow diagram illustrating a method for translating between gestures in accordance with some examples.

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

## 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. 7. 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 three-dimensional environments. Such techniques are described in the context of a computer system executing a 3D framework and a separate 2D framework. It should be recognized that other

configurations can be used with techniques described herein. For example, the 3D framework and/or the 2D frame can execute on a separate device using techniques described herein. In addition, techniques optionally complement or replace other techniques for managing three-dimensional 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/or performs different gestures, computer system 300 detects and processes the actions (e.g., movements and/or gestures of the user) 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 330. 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 of a scene graph (e.g., scene graph 400) in accordance with some examples. It should be recognized that the block diagram is not meant to be limiting and that such is used for discussion purposes. In some examples, scene graph 400 is a topological representation of a scene, with logical entities as nodes. In such examples, scene graph 400 can encode entities, their relationships, and operations required to render content. For example, world 410 of scene graph 400 can be a root node of scene graph 400 and include one or more render operations for the 3D environment.

**[0061]** From world 410, scene graph 400 can include one or more branch and/or leaf nodes that correspond to entities in the 3D environment and their respective rendering operations. For example, first entity 420 can correspond to visible content in the 3D environment and include operations to be performed to render the visible content. Such operations can include a position and/or orientation of the visible content along with textures to be used for the rendering process of the visible content. In some examples, the rendering process for the visible content is performed by 3D framework 320 via one or more APIs between 3D environment process 310 and 3D framework 320. Such operations can be performed without the use of 2D framework 330 as the visible content does not include any content rendered and/or managed by 2D framework 330.

**[0062]** In addition to first entity 420, scene graph 400 also includes first group 430 as a branch node of world 410. In some examples, first group 430 represents multiple entities that are related to each other. For example, second entity 440 can be a leaf node of first group 430 and represent more visible content that is rendered via 3D framework 320 without the use of 2D framework 330. In some examples, third entity 450 is another leaf node of first group 430. Third entity 450 can be a hidden entity that does not directly correspond to visible content. For example, third entity 450 might not by itself include any visible user-interface elements. Instead, third entity 450 can be a proxy for a 2D entity that is managed by 2D framework 330.

**[0063]** In some examples, the 2D entity includes information related to content rendered via 2D framework 330 (sometimes referred to as 2D content and/or 2D rendered content herein) and communicates such information and/or rendered texture to third entity 450 for use by 3D framework 320 when rendering. For example, a respective application (e.g., first user application 350 and/or second user application 360) can send a request to 3D environment process 310 to add a 2D object to the 3D environment. In response to the request, 3D environment process 310 can add a 3D object to a scene graph for the 3D environment, the 3D object is not visible in the 3D environment and is intended to hold a location and/or orientation for 2D content. In some examples, the 3D object can have a size corresponding to a size of the 2D content and be located at a position and/or an orientation requested by the respective application. In such examples, the 3D object can also include a mesh (e.g., a 3D mesh) that is used by 3D framework 320 when adding 2D rendered content. For example, 2D content is rendered by 2D framework 330 and placed on the 3D mesh by 3D framework 320 so that rendering the 3D environment by 3D framework 320 includes rendering the 2D rendered content within the 3D environment.

**[0064]** In some examples, third entity 450 is used to inform 3D framework 320 when to update. For example, 2D framework 330 and/or another process can detect an event that requires an update and, in response, send a request to 3D framework 320 via third entity 450 to update the 3D environment based on information associated with the 2D entity. In some examples, the combination of third entity 450 and the 2D entity allows 2D framework 330 to communicate with 3D framework 320, whereas typically the two frameworks would be completely independent.

**[0065]** In some examples, third entity 450 is attached to and/or configured relative to another 3D entity (e.g., second entity 440). In such examples, third entity 450 can ensure that a position and/or orientation of third entity 450 stays consistent with and/or maintains relative positioning and/or orientation with second entity 440 so that rendered content received by the 2D entity can move with second entity 440 without the 2D entity and/or 2D framework 330 needing to track where second entity 440 is located.

**[0066]** In some examples, third entity 450 includes information such as transformations (e.g., rotating, shrinking, enlarging, stretching, modifying a shape, and/or modifying a color characteristic) to be performed to content received from the 2D entity, collision dynamics to indicate how third entity 450 (and, by consequence to the 2D content tracking location and/or orientation of third entity 450, the 2D content) reacts to collisions with other 3D objects in the 3D environment, and/or physics dynamics to indicate how physics affects position and/or orientation of third entity 450 (and, by consequence to the 2D content tracking location and/or orientation of third entity 450, the 2D content) within the 3D environment.

**[0067]** While the above discussion is with respect to a single 3D environment, it should be recognized that 3D environment process 310 can manage multiple different environments (e.g., serially and/or simultaneously) and that different environments can include overlapping visible and/or hidden entities.

**[0068]** FIG. 5 is a block diagram illustrating a process (e.g., process 500) for processing gestures in a 3D environment according to some examples. Some operations in process 500 are, optionally, combined, the orders of some operations are, optionally, changed, and some operations are, optionally, omitted.

**[0069]** In some examples, process 500 is performed by a process (e.g., a system process or an application process) executing on a computer system, such as 3D environment process 310 described above. In such examples, the process can detect a 3D gesture and conditionally translate the 3D gestures into a 2D gesture.

**[0070]** Process 500 begins at 510, where the process detects a first 3D gesture. In some examples, the first 3D gesture includes 6 degrees of freedom, such as an x, y, and z value. For example, the first 3D gesture can correspond to movement of a user (e.g., a hand, an eye, and/or a leg) and be captured via one or more cameras in communication with the computer system. Other examples of input devices that can be used include a camera, a motion sensor,

a depth sensor, a remote control, a gyroscope, an accelerometer, a touch-sensitive surface, and/or a physical input mechanism (e.g., a keyboard, a mouse, a rotatable input mechanism, and/or a physical button).

**[0071]** At 520, in response to detecting the first 3D gesture, the process determines which entity in a 3D environment (e.g., a virtual or a physical 3D environment) that the first 3D gesture is directed to. In some examples, the entity is determined using a scene graph (e.g., scene graph 400 of FIG. 4). For example, the process can determine that the first 3D gesture is directed to a particular location within the 3D environment and, based on the particular location and the scene graph, determines that the first 3D gesture is directed to a particular entity in the scene graph. In some examples, the particular entity is a visible entity (e.g., content that corresponds to a 3D framework (e.g., 3D framework 320) and not a 2D framework (e.g., 2D framework 330)). In other examples, the particular entity is a hidden entity (e.g., content that corresponds to the 2D framework and not the 3D framework).

**[0072]** At 530, in response to determining that the first 3D gesture is directed to a visible entity in the scene graph, the first process sends an indication of the first 3D gesture to the 3D framework and/or an application (e.g., a user application executing on the computer system, such as first user application 350 and/or second user application 360) corresponding to the visible entity. In some examples, the indication of the first 3D gesture is sent to the application corresponding to the visible entity via the 3D framework. In some examples, the indication of the first 3D gesture includes and/or is an indication of a type of gesture and does not specify one or more locations and/or orientations of an input that is detected (e.g., the indication is not at the level of joints and positions of fingers but rather at the level of communicating a category).

**[0073]** At 540, in response to determining that the first 3D gesture is directed to a hidden entity in the scene graph, the first process determines whether the 2D framework is aware of a type of gesture corresponding to the first 3D gesture. For example, the 2D framework can be configured to recognize some types of 3D gestures and/or some 3D gestures are not specific to a 3D coordinate space and therefore do not need to be translated to a type of 2D gesture.

**[0074]** At 550, in response to determining that the 2D framework is aware of a type of gesture corresponding to the first 3D gesture, the first process sends the indication of the first

3D gesture to the 2D framework and/or an application (e.g., a user application executing on the computer system, such as first user application 350 and/or second user application 360) corresponding to the hidden entity. In some examples, the indication of the first 3D gesture is sent to the application corresponding to the hidden entity via the 2D framework. In some examples, the indication of the first 3D gesture includes and/or is an indication of a type of gesture and does not specify one or more locations and/or orientations of an input that is detected (e.g., the indication is not at the level of joints and positions of fingers but rather at the level of communicating a category).

**[0075]** At 560, in response to determining that the 2D framework is not aware of a type of gesture corresponding to the first 3D gesture, the first process translates the first 3D gesture into a first 2D gesture (e.g., a gesture that is a type of 2D gesture) and sends an indication of the first 2D gesture to the 2D framework and/or the application corresponding to the hidden entity. In some examples, translating the first 3D gesture includes identifying a 2D gesture corresponding to the first 3D gesture. In some examples, translating the first 3D gesture includes removing data from the first 3D gesture to create the first 2D gesture. In some examples, translating the first 3D gesture includes removing an axis (e.g., removing z axis from x, y, and z) and/or reducing from 6 degrees of freedom to 2 degrees of freedom. It should be recognized that the examples for translating the first 3D gesture described above are not an exhaustive list and that some examples can be combined with others. In some examples, the indication of the first 2D gesture is sent to the application corresponding to the hidden entity via the 2D framework. In some examples, the indication of the first 2D gesture includes and/or is an indication of a type of gesture and does not specify one or more locations and/or orientations of an input that is detected (e.g., the indication is not at the level of joints and positions of fingers but rather at the level of communicating a category).

**[0076]** In some examples, an indication that is sent can correspond to a first type of coordinate space (e.g., a world space (e.g., a specific location in the 3D environment) and/or a user space (e.g., a location corresponding to a field of view of the computer system)). In such examples, after sending an indication of a respective gesture, the process receives, from an application, a request for a second type of coordinate space different from the first type of coordinate space. In response to the request for the second type of coordinate space, first process can convert the indication that was sent to a different coordinate space and send the converted indication to the application. In other examples, before the first 3D gesture is

detected, the application defines a type of coordinate space for the process to use when communicating with the computer system.

**[0077]** In some examples, an indication that is sent can correspond to a first type of unit (e.g., a type of measurement). In such examples, after sending an indication of a respective gesture, the process receives, from an application, a request for a second type of unit different from the first type of unit. In response to the request for the second type of unit, first process can convert the indication that was sent to the second type of unit and send the converted indication to the application. In other examples, before the first 3D gesture is detected, the application defines a type of unit for the process to use when communicating with the computer system.

**[0078]** While the discussion above is with respect to translating from a 3D gesture to a 2D gesture, it should be recognized that techniques described herein can be used for translating from a 2D gesture to a 3D gesture depending on a destination. For example, a 2D gesture can be detected and, in accordance with a determination that the 2D gesture is being sent to a component requesting and/or needing a 3D gesture, the first process can translate the 2D gesture to a first 3D gesture using the opposite of one or more of the techniques described above with respect to translating from a 3D gesture to a 2D gesture. In some examples, translation from a 2D gesture to a 3D gesture can be used in the context of a 2D window and/or application providing the first process a first 2D gesture and needing a 3D gesture corresponding to the first 2D gesture so that the application can send the 3D gesture to a 3D framework.

**[0079]** FIG. 6 is a flow diagram illustrating a method (e.g., method 600) for integrating a 2D framework with a 3D framework in accordance with some examples. Some operations in method 600 are, optionally, combined, the orders of some operations are, optionally, changed, and some operations are, optionally, omitted. In some examples, method 600 is performed by a computer system (e.g., 300).

**[0080]** At 602, the computer system receives, from an application (e.g., 340 and/or 350) (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 (e.g., a device, a personal device, a user device, and/or a head-mounted display (HMD))), a request to add (e.g., display, configure, and/or place) a two-dimensional (2D) entity (e.g., the 2D entity corresponding to

450) (e.g., a 2D object, a 2D model, a 2D user interface, a 2D window, and/or a 2D user interface element) at a first location in a three-dimensional (3D) environment (e.g., a virtual reality environment, a mixed reality environment, and/or an augmented reality environment), wherein the 3D environment includes one or more 3D entities (e.g., 420, 440, and/or 450) (e.g., a 3D object, a 3D model, a 3D user interface, a 3D window, and/or a 3D user-interface element). 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. 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). In some examples, receiving the request to add the 2D entity at the first location in the 3D environment includes detecting, via one or more input devices in communication with the computer system, 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). In some examples, one or more operations of method 600 are performed via a process (e.g., 310) of the computer system.

**[0081]** At 604, in response to receiving the request to add the 2D entity at the first location in the 3D environment, the computer system adds a first 3D entity (e.g., 450) (e.g., a hidden or proxy entity) to the first location in the 3D environment. In some examples, adding the first 3D entity does not include displaying the first 3D entity. In some examples, the first 3D entity is not visible in the 3D environment. In some examples, the first 3D entity is different and/or separate from the 2D entity.

**[0082]** At 606, in response to receiving the request to add the 2D entity at the first location in the 3D environment, the computer system renders (e.g., synthesizes, generates, and/or creates), via a 2D framework (e.g., 330) (e.g., of the computer system or in

communication with the computer system) (e.g., a 2D user interface framework) (e.g., and not a 3D framework (e.g., 320)) (e.g., software that includes one or more predefined user-interface elements and/or one or more operations to build, generate, render, enable interactions with, and/or display a user interface and/or user interface element in two dimensions (e.g., and not three dimensions)), a representation (e.g., a graphical and/or visual representation) of the 2D entity.

**[0083]** At 608, in response to receiving the request to add the 2D entity at the first location in the 3D environment, the computer system renders, via a 3D framework (e.g., 320) (e.g., of the computer system or in communication with the computer system) (e.g., software that includes one or more predefined user-interface elements and/or one or more operations to build, generate, render, enable interactions with, and/or display a user interface and/or user interface element in three dimensions (e.g., and not two dimensions)), a representation (e.g., a graphical and/or visual representation) of a second 3D entity (e.g., 450) (e.g., different and/or separate from the first 3D entity) by (e.g., rendering the representation of the second 3D entity includes) performing (e.g., via the 3D framework) one or more operations on the representation of the 2D entity. In some examples, the second 3D entity is a visible representation of the first 3D entity. In some examples, the second 3D entity is different and/or separate from the 2D entity.

**[0084]** At 610, in response to receiving the request to add the 2D entity at the first location in the 3D environment, the computer system renders, via the 3D framework, the representation of the second 3D entity by placing (e.g., via the 3D framework) the representation of the second 3D entity at the first location (e.g., at a location corresponding to the first 3D entity). In some examples, after placing the representation of the second 3D entity at the first location, the computer system displays, via a display generation component in communication with the computer system, the 3D environment including display of at least a portion of the representation of the second 3D entity at the first location.

**[0085]** In some examples, an observable (e.g., perceivable, visible, and/or audible) representation (e.g., a visual and/or an auditory representation) of the first 3D entity is not present (e.g., visible, displayed, and/or output) in the 3D environment (e.g., the first 3D entity is used for tracking purposes for the application, the 3D framework, and/or the 2D framework, such as to track position, orientation, and/or pose for the second 3D entity and/or information related to the 2D entity).

**[0086]** In some examples, performing the one or more operations include rotating the representation of the 2D entity, shirking the representation of the 2D entity, enlarging the representation of the 2D entity, stretching the representation of the 2D entity, modifying a shape of the representation of the 2D entity, modifying a color characteristic (e.g., hue, saturation, and/or brightness) of the representation of the 2D entity, or one or more combinations thereof (e.g., based on information (e.g., position, orientation, and/or pose) included in the first 3D entity) (e.g., based on information (e.g., a size and/or a color characteristic) included in the second 2D entity) (e.g., based on information (e.g., a color characteristic and/or a physics setting to augment an appearance of the representation of the 2D entity) corresponding to the 3D environment).

**[0087]** In some examples, the computer system sends, to a display process (e.g., 340) (e.g., of the computer system and/or of a display generation component (e.g., of a head-mounted display (HMD) device) in communication with the computer system) (e.g., a system process or another application different from the application), a request to display a representation of the one or more 3D entities (e.g., the computer system causes the display process to display the representation of the one or more 3D entities). In some examples, the display process is different from the process of the computer system.

**[0088]** In some examples, after placing the representation of the second 3D entity at the first location, the computer system detects movement of a visible representation (e.g., visible within the 3D environment) of a third 3D entity of the one or more 3D entities from a second location to a third location. In some examples, the third location is different from the second location. In some examples, the second location is a first distance from the first location (e.g., zero or more units). In some examples, the third location is different from the first location. In some examples, the visible representation of the third 3D entity is not currently visible as a result of a current orientation of a view area but is visible as a result of another orientation different from the current orientation of the view area. In some examples, in response to detecting the movement of the visible representation of the third 3D entity, the computer system places (e.g., via the 3D framework) the representation of the second 3D entity at a fourth location different from the first location. In some examples, the fourth location is the first distance from the third location (e.g., the representation of the second 3D entity maintains a relative position to the visible representation of the third 3D entity).

**[0089]** In some examples, the computer system receives (e.g., via the 2D framework and/or the application) a change to the 2D entity. In some examples, the change is caused by the application. In some examples, the change is caused by an interaction of the representation of the 2D entity with a representation of a 3D entity in the 3D environment. In some examples, the change is caused by 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) detected via one or more input devices (e.g., a camera (e.g., a telephoto camera, a wide-angle camera, and/or an ultra-wide-angle camera), a microphone, a sensor (such as a heart rate sensor), a touch-sensitive surface, a mouse, a keyboard, a touch pad, and/or an input mechanism (e.g., a physical input mechanism, such as a rotatable input mechanism and/or a button)) in communication with the computer system. In some examples, in response to receiving the change to the 2D entity, the computer system renders, via the 3D framework, a second representation of the 2D entity, wherein the second representation is different from (e.g., a different appearance, location, orientation, and/or pose) the representation of the 2D entity.

**[0090]** In some examples, the computer system detects, via one or more input devices (e.g., a camera (e.g., a telephoto camera, a wide-angle camera, and/or an ultra-wide-angle camera), a microphone, a sensor (such as a heart rate sensor), a touch-sensitive surface, a mouse, a keyboard, a touch pad, and/or an input mechanism (e.g., a physical input mechanism, such as a rotatable input mechanism and/or a button)) (e.g., in communication with the computer system), 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) (e.g., a 3D gesture) directed to the representation of the second 3D entity. In some examples, in response to detecting the input directed to the representation of the second 3D entity and in accordance with a determination that a first set of one or more criteria (e.g., as described below with respect to method 700) is satisfied, the computer system translates the input to a second input (e.g., generating the second input) (e.g., a 2D representation of the input) different from the

input (e.g., as described below with respect to method 700). In some examples, the first set of one or more criteria includes a criterion that is satisfied when the input is a first type of input (e.g., as described below with respect to method 700). In some examples, in response to detecting the input directed to the representation of the second 3D entity and in accordance with a determination that a first set of one or more criteria (e.g., as described below with respect to method 700) is satisfied, the computer system sends, to the application (e.g., via the 2D framework), an indication of the second input (e.g., without sending an indication of the input). In some examples, in response to detecting the input directed to the representation of the second 3D entity and in accordance with a determination that a second set of one or more criteria is satisfied (e.g., as described below with respect to method 700), the computer system (and/or the process of the computer system) sends, to the application (e.g., via the 2D framework), an indication of the input (e.g., without sending the indication of the second input). 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, the second set of one or more criteria includes a criterion that is satisfied when the first set of one or more criteria is not satisfied. In some examples, in response to detecting the input directed to the representation of the second 3D entity and in accordance with a determination that a third set of one or more criteria (e.g., as described below with respect to method 700) is satisfied, the computer system (and/or the process of the computer system) sends, to the application (e.g., via the 2D framework), the indication of the second input and the indication of the input. In some examples, the third set of one or more criteria is different from the first set of one or more criteria and/or the second set of one or more criteria. In some examples, the third set of one or more criteria includes a criterion that is satisfied when the first set of one or more criteria and/or the second set of one or more criteria is not satisfied.

**[0091]** In some examples, the computer system detects an interaction (e.g., a collision and/or an effect of a proximity (e.g., gravity, a pull effect, and/or a push effect)) between the representation of the second 3D entity and a representation of a fourth 3D entity of the one or more 3D entities. In some examples, in response to detecting the interaction, the computer system renders (e.g., synthesizes, generates, and/or creates), via the 2D framework, an updated representation (e.g., a graphical and/or visual representation) of the 2D entity. In some examples, the updated representation of the 2D entity is different from the representation of the 2D entity. In some examples, in response to detecting the interaction, the computer system renders, via the 3D framework, an updated representation (e.g., a

graphical and/or visual representation) of the second 3D entity. In some examples, the updated representation of the second 3D entity is different from the representation of the second 3D entity.

**[0092]** In some examples, the computer system receives, from the application, information (e.g., an effect of the 2D entity on one or more other representations in the 3D environment, such as a glow or other outward effect on an area outside of a representation of the 2D entity) corresponding to the 2D entity. In some examples, the information is included in the request to add the 2D entity at the first location in the 3D environment. In some examples, after receiving the request to add the 2D entity at the first location in the 3D environment, the computer system renders (e.g., via the 3D framework) the 3D environment (e.g., renders one or more representations of one or more visible entities in the 3D environment) (e.g., based on the information corresponding to the 2D entity), wherein the information affects rendering a visual representation of a fifth 3D entity in the 3D environment, and wherein the fifth 3D entity does not correspond to the 2D entity. In some examples, the information affects rendering the visual representation of the fifth 3D entity in the 3D environment by causing a color of the fifth 3D entity to be based on a color of the 2D entity.

**[0093]** In some examples, the first 3D entity is added to a graph (e.g., 400) (e.g., a scene graph) of the 3D environment.

**[0094]** In some examples, the first 3D entity is included in a second 3D environment different from the 3D environment. In some examples, the second 3D environment is not displayed concurrently with the 3D environment. In some examples, the second 3D environment is displayed in response to a request to display the second 3D environment (e.g., and not the 3D environment). In some examples, updates to the first 3D entity affects the 3D environment and the second 3D environment.

**[0095]** In some examples, the first 3D entity includes a size (e.g., corresponding to the 2D entity, the first 3D entity, the second 3D entity, and/or the representation of the second 3D entity), a location (e.g., corresponding to the 2D entity, the first 3D entity, the second 3D entity, and/or the representation of the second 3D entity) (e.g., within the 3D environment), an orientation (e.g., corresponding to the 2D entity, the first 3D entity, the second 3D entity, and/or the representation of the second 3D entity) (e.g., within the 3D environment), a pose

(e.g., corresponding to the 2D entity, the first 3D entity, the second 3D entity, and/or the representation of the second 3D entity) (e.g., within the 3D environment), or any combination thereof.

**[0096]** In some examples, the computer system receives, from a second application (e.g., 350 and/or 360) (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 a computer system (e.g., a device, a personal device, a user device, and/or a head-mounted display (HMD))) different from the application, a request to add a sixth 3D entity (e.g., corresponding to another 2D entity and/or not corresponding to another 2D entity) to the 3D environment. In some examples, the second application is not associated with the application. In some examples, the second application is a different type of application than the application. In some examples, after receiving the request to add the sixth 3D entity to the 3D environment, the computer system renders, via the 3D framework, a representation of the sixth 3D entity. In some examples, the computer system causes concurrent display, in the 3D environment, of the representation of the sixth 3D entity and the representation of the third 3D entity.

**[0097]** In some examples, the application causes the one or more 3D entities to be rendered for (and/or displayed in) the 3D environment.

**[0098]** Note that details of the processes described above with respect to method 600 (e.g., FIG. 6) 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 600. For example, the 3D environment of method 700 is the 3D environment of method 600. For brevity, these details are not repeated below.

**[0099]** FIG. 7 is a flow diagram illustrating a method (e.g., method 700) for translating between gestures 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.

**[0100]** In some examples, method 700 is performed at a computer system (e.g., 300) (e.g., a device, a personal device, a user device, and/or a head-mounted display (HMD)) in communication with one or more input devices (e.g., a camera, a touch-sensitive surface, a motion detector, and/or a microphone). 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. In some examples, method 700 is performed by a daemon (e.g., 310), a system process (e.g., 310), and/or a user process (e.g., 310, 350, and/or 360).

**[0101]** At 702, the computer system detects, via the one or more input devices, a first input (e.g., 510) (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 a three-dimensional (3D) environment (e.g., a virtual reality environment, a mixed reality environment, and/or an augmented reality environment).

**[0102]** At 704, in response to detecting the first input corresponding to the 3D environment and in accordance with a determination (e.g., 520) that a first set of one or more criteria is satisfied, wherein the first set of one or more criteria includes a criterion that is satisfied when the input is directed to a first entity (e.g., 450) (e.g., an object, a model, a user interface, a window, an element in a scene graph, and/or a user-interface element) of a first type (e.g., a hidden or proxy entity) in the 3D environment, the computer system translates (e.g., 560) the first input to a second input different from the first input. In some examples, the second input is a different format than the first input. In some examples, the second input corresponds to the first input.

**[0103]** At 706, in response to detecting the first input corresponding to the 3D environment and in accordance with a determination that the first set of one or more criteria is satisfied, the computer system sends (e.g., 560), to a first application (e.g., 350 and/or 360) (e.g., a user application and/or an application installed on the computer system (e.g., by a user)) (e.g., of the computer system (e.g., corresponding to the first entity)) (e.g., via a first framework (e.g., 330) (e.g., of the computer system or in communication with the computer system) (e.g., a two-dimensional (2D) framework and/or a 2D user interface framework) (e.g., and not a 3D framework) (e.g., software that includes one or more predefined user-interface elements and/or one or more operations to build, generate, render, enable interactions with, and/or display a user interface and/or user interface element in two dimensions (e.g., and not three dimensions))), an indication of the second input (e.g., without

sending an indication of the first input). In some examples, the first application corresponds to the first entity.

**[0104]** At 708, in response to detecting the first input corresponding to the 3D environment and in accordance with a determination that a second set of one or more criteria is satisfied, wherein the second set of one or more criteria includes a criterion that is satisfied when the input is directed to a second entity of a second type (e.g., 420 and/or 440) (e.g., a 3D object, a 3D entity, a 3D model, a 3D user interface, a 3D window, and/or a 3D user-interface element) in the 3D environment, the computer system sends (e.g., 430), to a second application (e.g., of the computer system) (e.g., via a second framework (e.g., 320) different from the first framework (e.g., of the computer system or in communication with the computer system) (e.g., a three-dimensional (3D) framework and/or a 3D user interface framework) (e.g., and not a 2D framework) (e.g., software that includes one or more predefined user-interface elements and/or one or more operations to build, generate, render, enable interactions with, and/or display a user interface and/or user interface element in three dimensions (e.g., and not two dimensions))), an indication of the first input (e.g., without sending the indication of the second input), wherein the second type of entity is different from the first type of entity, and wherein the second set of one or more criteria is different from the first set of one or more criteria. In some examples, the second application corresponds to the second entity. In some examples, a method includes: displaying 2D content; while displaying the 2D content, detecting, via one or more input devices, a 3D gesture directed to the 2D content; in response to detecting the 3D gesture directed to the 2D content, translating the 3D gesture to a 2D gesture and providing the 2D gesture to an application (e.g., via a 2D framework as described above) corresponding to the 2D content (e.g., the application caused display of and/or includes the 2D content).

**[0105]** In some examples, an observable (e.g., perceivable, visible, and/or audible) representation (e.g., a visual and/or an auditory representation) of the first entity of the first type (e.g., a hidden entity corresponding to a 2D entity, as described above with respect to method 600) is not present (e.g., visible, displayed, and/or output) in the 3D environment (e.g., the first entity of the first type is used for tracking purposes for the first application, the 3D framework, and/or the 2D framework, such as to track position, orientation, and/or pose for a representation of an entity related to, corresponding to, and/or associated with the first entity of the first type).

**[0106]** In some examples, an observable (e.g., perceivable, visible, and/or audible) representation (e.g., a visual and/or an auditory representation) of the first entity of the second type (e.g., a visible entity, as described above with respect to method 600) is present (e.g., visible, displayed, and/or output) in the 3D environment.

**[0107]** In some examples, in response to detecting the first input corresponding to the 3D environment and in accordance with a determination that a third set of one or more criteria is satisfied, wherein the third set of one or more criteria includes a criterion that is satisfied when the input is a first type of input (e.g., a type of gesture known by the first application and/or the 2D framework) directed to the first entity of the first type in the 3D environment, the computer system sends (e.g., 550), to the first application (e.g., via the first framework), the indication of the first input (e.g., without translating the first input to the second input), wherein the third set of one or more criteria is different from the first set of one or more criteria and the second set of one or more criteria, and wherein the first set of one or more criteria includes a criterion that is satisfied when the input is a second type of input (e.g., a type of gesture not known by the first application and/or the 2D framework) different from the first type of input.

**[0108]** In some examples, the first input includes a 3D gesture (e.g., a hand pose in three dimensions and/or movement in three dimensions). In some examples, the first input is the 3D gesture.

**[0109]** In some examples, the first input includes an air gesture. In some examples, the first input is the air gesture.

**[0110]** In some examples, the first input includes a gaze of a user directed to a location (e.g., a location corresponding to the first entity and/or the second entity) in the 3D environment. In some examples, the first input is the gaze of the user directed to the location in the 3D environment.

**[0111]** In some examples, in response to detecting the first input corresponding to the 3D environment and in accordance with a determination that a fourth set of one or more criteria is satisfied, wherein the fourth set of one or more criteria includes a criterion that is satisfied when the input is directed to a third entity (e.g., the first entity, the second entity, or another entity different from the first entity and the second entity) of a third type in the 3D environment, the computer system sends, to a fourth application (e.g., via the first framework

and/or the second framework), an indication of a location corresponding to a first type of coordinate space for the first input, an orientation corresponding to the first type of coordinate space for the first input, a pose corresponding to the first type of coordinate space for the first input, a magnitude corresponding to the first type of coordinate space for the first input, or any combination thereof, wherein the first type of coordinate space is defined by the fourth application. In some examples, the fourth application is the first application or the second application. In some examples, the indication corresponding to the first type of coordinate space is sent with an indication of an input (e.g., the indication of the first input and/or the indication of the second input). In some examples, the first type of coordinate space is a coordinate space with respect to a window of the fourth application (e.g., application centric, such as a distance from the window) (e.g., and not with respect to a location in the 3D environment other than a location corresponding to the window). In some examples, the indication corresponding to the first type of coordinate space includes an identification of a type of gesture. In response to detecting the first input corresponding to the 3D environment and in accordance with a determination that a fifth set of one or more criteria is satisfied, wherein the fifth set of one or more criteria includes a criterion that is satisfied when the input is directed to a fourth entity of a fourth type in the 3D environment, the computer system sends, to the fourth application (e.g., via the first framework and/or the second framework), an indication of a location corresponding to a second type of coordinate space for the first input, an orientation corresponding to the second type of coordinate space for the first input, a pose corresponding to the second type of coordinate space for the first input, a magnitude corresponding to the second type of coordinate space for the first input, or any combination thereof, wherein the second type of coordinate space is defined by the fourth application, wherein the fifth set of one or more criteria is different from the fourth set of one or more criteria, wherein the fourth type of entity is different from the third type of entity, and wherein the second type of coordinate space is different from the first type of coordinate space (e.g., the indication corresponding to the second type of coordinate space is different from the indication corresponding to the first type of coordinate space). In some examples, the fourth entity is different from the third entity. In some examples, the third type of entity is different from the first type of entity and the second type of entity. In some examples, the fourth type of entity is different from the first type of entity and the second type of entity. In some examples, the indication corresponding to the second type of coordinate space is sent with an indication of an input (e.g., the indication of the first input and/or the indication of the second input). In some examples, the second type of coordinate space is a coordinate space

with respect to the 3D environment (e.g., a location within the 3D environment, such as a location of a user, a location of a view point, and/or a location of an object other than the window) (e.g., world centric, such as a distance from a user, and/or world centric, such as a stance from a location in the world) (e.g., and not with respect to a location corresponding to the window). In some examples, the indication corresponding to the first type of coordinate space and the indication corresponding to the second type of coordinate space is sent to the fourth application (e.g., the fourth application is defined to receive indications corresponding to the first type of coordinate space and the second type of coordinate space). In some examples, when the fourth application does not include a definition of a type of coordinate space, the computer system sends, to the fourth application, the indication corresponding to the first type of coordinate space and the indication corresponding to the second type of coordinate space. In some examples, the computer system receives, from the fourth application, a request for an indication for an input corresponding to a particular type of coordinate space. In some examples, in response to receiving the request for the indication for the input corresponding to the particular type of coordinate space, the computer system sends, to the fourth application, the indication for the input corresponding to the particular type of coordinate space. In some examples, the indication corresponding to the second type of coordinate space includes an identification of a type of gesture.

**[0112]** In some examples, the second application is the first application. In some examples, the second application is different from the first application.

**[0113]** In some examples, the second input includes a two-dimensional representation of the first input (e.g., a location, orientation, and/or pose with two or less dimensions).

**[0114]** In some examples, translating the first input to the second input includes modifying a representation of a respective input from having six degrees of freedom to two degrees of freedom.

**[0115]** 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 600 optionally includes one or more of the characteristics of the various methods described above with reference to method 700. For example, the first entity of the first type of method 700 is the first 3D entity of method 600. For brevity, these details are not repeated below.

[0116] In some examples, a view for hosting a 3D framework simulation within a 2D framework view hierarchy is provided. In some examples, the view is defined in a cross-import overlay of a 3D framework (e.g., a 3D UI framework) and a 2D framework (e.g., a 2D UI framework). In some examples, the view can be implemented as a wrapper for a view of the 3D framework. In other examples and/or other types of devices, the view can be implemented as an integration of the 3D framework with the 2D framework via one or more private serial peripheral interfaces (SPIs).

[0117] In some examples, the view uses inline closure syntax to bridge between the declarative world of the 2D framework and the imperative world of the 3D framework. In some examples, within these closures, the view provides a mutable `Content` struct. This struct serves as a container for entities in the view's hierarchy, as well as an entry point to one or more top-level APIs corresponding to the 3D framework.

[0118]

[0119] Here is a simple example of how one could use the view:

[0120]

[0121] struct MyCoolSimulation: View {

[0122]     @State var state = SimulationState()

[0123]     var body: some View {

[0124]         View { content in

[0125]             content.add(state.modelEntity)

[0126]             // subscribe to scene events, do other setup as needed

[0127]         }

[0128]     }

[0129] }

[0130]

[0131]     @Observable

[0132]     class SimulationState {

[0133]         var modelEntity = ModelEntity(mesh: .generateSphere(radius: 0.1))

[0134]     }

[0135]     ``

[0136]

[0137]     Handling state updates for the view can be done by adding an `update` closure, evaluated any time the containing view's body is re-evaluated:

[0138]

[0139]     struct MyCoolSimulation: View {

[0140]         @State var state = SimulationState()

[0141]         var showModel: Bool

[0142]

[0143]         var body: some View {

[0144]             View { content in

[0145]                 content.add(state.modelEntity)

[0146]             } update: { \_ in

[0147]                 state.modelEntity.isEnabled = showModel

[0148]             }

[0149]         }

[0150]     }

[0151]     ``

[0152]

[0153] In some examples, note that the view omits the use of a context or coordinator, as developers can instead use standard constructs corresponding to the 2D framework like `@State` and `@Environment` as needed.

[0154] In the prior examples, we generated a simple mesh synchronously for use in our view. In some examples, a developer working with the 3D framework will instead load larger models from disk or a URL, which do not need to be done synchronously on the main actor. In some examples, the view provides a simple convenience for this use case, as its primary `make` closure is actually `async`:

[0155]

[0156] `struct MyCoolSimulation: View {`

[0157] `var body: some View {`

[0158] `View { content in`

[0159] `if let robot = try? await ModelEntity(named: "robot") {`

[0160] `content.add(robot)`

[0161] `}`

[0162] `}`

[0163] `}`

[0164] `}`

[0165] ````

[0166]

[0167] In some examples, until the `make` closure has completed, the `view` is not considered visible, and a placeholder view is shown instead. In some examples, a developer can customize this placeholder view using the optional `placeholder` `ViewBuilder`, such as

to show a `ProgressView` spinner. In some examples, a developer can use another instance of the view to display a 3D placeholder model:

[0168]

[0169]     struct MyCoolSimulation: View {

[0170]         @State var state = SimulationState()

[0171]         var body: some View {

[0172]             View { content in

[0173]                 if let robot = try? await ModelEntity(named: "robot") {

[0174]                     content.add(robot)

[0175]                 }

[0176]             } placeholder: {

[0177]                 View { content in

[0178]                     content.add(state.placeholderSphere)

[0179]                 }

[0180]             }

[0181]         }

[0182]     }

[0183]     ``

[0184]

[0185]     In some examples, the `Content` struct provided to the view's closures is a generic type conforming to a protocol. In such examples, the protocol represents the core interface common to some permutations of the view. Concrete types conforming to this protocol can add additional functionality specific to their configuration. For example, an isolated view

could provide additional functionality in its `Content` struct that a shared view would not have access to (more discussion on isolation later).

[0186]

[0187] In some examples, the protocol itself is defined like so:

[0188]

[0189] @available(mixedrealityOS, phoneOS, computerOS)

[0190] public protocol Protocol {

[0191] associatedtype Entities: EntityCollection

[0192]

[0193] var entities: Entities { get nonmutating set }

[0194]

[0195] func subscribe<E: Event>(to event: E.Type,

[0196] on sourceObject: EventSource?,

[0197] componentType: Component.Type?,

[0198] \_handler: @escaping (E) -> Void) -> EventSubscription

[0199] }

[0200] ```

[0201]

[0202] In some examples, processes add and/or remove entities using convenience APIs via the protocol:

[0203]

[0204] @available(mixedrealityOS, phoneOS, computerOS)

```
[0205]     extension Protocol {  
[0206]         public func add(_ entity: Entity)  
[0207]         public func remove(_ entity: Entity)  
[0208]     }  
[0209]     ``````  
[0210]  
[0211]  
[0212]     struct MyCoolSimulation: View {  
[0213]         var body: some View {  
[0214]             View { content in  
[0215]                 await addPrimaryContent(to: &content)  
[0216]                 Task { [content] in  
[0217]                     // We can still add this secondary model within the scope of  
[0218]                     // this Task!  
[0219]                     if let secondaryModel = try? await ModelEntity(named: "...") {  
[0220]                         content.add(secondaryModel)  
[0221]                     }  
[0222]                 }  
[0223]             }  
[0224]         }  
[0225]     }
```

[0226]     '''

[0227]

[0228]

[0229]     @available(mixedrealityOS, phoneOS, computerOS)

[0230]     extension Entity.ChildCollection: EntityCollection {

[0231]

[0232]     }

[0233]     '''

[0234]

[0235]     In some examples, to subscribe to events, like collision events between entities, clients can use the `subscribe(to:on:componentType:)` method on the protocol, or one of the convenience methods provided in an extension:

[0236]

[0237]     @available(mixedrealityOS, phoneOS, computerOS)

[0238]     extension Protocol {

[0239]         public func subscribe<E: Event>(to event: E.Type,

[0240]                                 on sourceObject: EventSource? = nil

[0241]                                 \_ handler: @escaping (E) -> Void) -> EventSubscription

[0242]

[0243]         public func subscribe<E: Event>(to event: E.Type,

[0244]                                 componentType: Component.Type?,

[0245]                                 \_ handler: @escaping (E) -> Void) -> EventSubscription

[0246] }

[0247]

[0248] @available(mixedrealityOS, phoneOS, computerOS)

[0249] public struct EventSubscription {

[0250] public func cancel()

[0251] }

[0252] ```

[0253]

[0254] In some examples, `EventSubscription` is defined in the 3D framework.

[0255] In some examples, events published through `EventSubscription` are scoped only to the entities contained within that view. In some examples, the event subscription will automatically terminate when the `EventSubscription` is destroyed. In some examples, the event subscription can be terminated early by calling `EventSubscription.cancel()`. In some examples, an EventSubscription will automatically cancel the associated event subscription on destruction.

[0256] For example, to subscribe to collision events between any two entities in a view:

[0257]

[0258] struct JarOfMarbles: View {

[0259] @State var state = SimulationState()

[0260] @State var subscription: EventSubscription?

[0261] var body: some View {

[0262] View { content in

[0263] let marbles = await state.loadMarbles()

```

[0264]         content.entities += marbles

[0265]         subscription = content.subscribe(to: CollisionEvents.Began.self) {

[0266]             state.playMarbleClinkSound(at: collision.position)

[0267]         }

[0268]     }

[0269] }

[0270] }

[0271] ```

[0272]

```

[0273] In some examples, the view provides content in a shared scene on mixed reality operating system, and content isolated to a 2D camera projection (AR or non-AR) on a phone operating system and/or a computer operating system. In such examples, these are defined like so:

```

[0274]

[0275] @available(mixedrealityOS)

[0276] @available(phoneOS)

[0277] @available(computerOS)

[0278] public struct Content: Protocol {

[0279]     ...

[0280] }

[0281]

[0282] @available(phoneOS, computerOS)

[0283] public struct Content: Protocol {

```

[0284] ...

[0285] }

[0286] ```

[0287]

[0288] In some examples, additional APIs can be provided on each type separately. In some examples, to convert coordinates between the 2D framework space and the 3D framework entity space, we introduce new methods on `Content` for a mixed reality operating system:

[0289]

[0290] @available(mixedrealityOS)

[0291] extension Content {

[0292] public func convert(

[0293] \_ point: SIMD3<Float>, from entity: Entity? = nil,

[0294] to space: CoordinateSpace

[0295] ) -> Point3D

[0296] public func convert(

[0297] \_ point: Point3D, from space: CoordinateSpace, to entity: Entity? = nil

[0298] ) -> SIMD3<Float>

[0299]

[0300] public func convert(

[0301] \_ rect: BoundingBox, from entity: Entity? = nil,

[0302] to space: CoordinateSpace

[0303] ) -> Rect3D

```

[0304]     public func convert(
[0305]         _ rect: Rect3D, from space: CoordinateSpace, to entity: Entity? = nil
[0306]     ) -> BoundingBox
[0307]
[0308]     public func convert(
[0309]         _ transform: AffineTransform3D, from space: CoordinateSpace,
[0310]         to entity: Entity? = nil
[0311]     ) -> Transform
[0312]     public func convert(
[0313]         _ transform: Transform, from entity: Entity? = nil,
[0314]         to space: CoordinateSpace
[0315]     ) -> AffineTransform3D
[0316] }
[0317] ```

```

[0318] In some examples, in this construction, the "to / from" `CoordinateSpace` represents the 2D framework reference space, while the other is implicitly the local space of the Content. In some examples, the 2D framework space is represented in points, with Y pointing down and a top-left-back origin, while the 3D framework space is represented in meters, with Y pointing up and a center origin. We also use the currency types for each framework (`SIMD3<Float>`, `BoundingBox`, and `Transform` as used in the 3D framework, as opposed to the Spatial types `Point3D`, `Rect3D`, and `AffineTransform3D` used in the 2D framework).

```

[0319] @available(phoneOS, computerOS)

```

```

[0320] extension Content {

```

[0321]       public func project(  
[0322]            vector: SIMD3<Float>, to space: CoordinateSpace  
[0323]       ) -> CGPoint?  
[0324]  
[0325]       public func unproject(  
[0326]            point: CGPoint, from space: CoordinateSpace, ontoPlane: float4x4,  
[0327]            relativeToCamera: Bool = false  
[0328]       ) -> SIMD3<Float>?  
[0329]       public func unproject(  
[0330]            point: CGPoint, viewport: CGRect  
[0331]       ) -> SIMD3<Float>?  
[0332]  
[0333]       public func ray(  
[0334]            through: CGPoint, in space: CoordinateSpace  
[0335]       ) -> (origin: SIMD3<Float>, direction: SIMD3<Float>)?  
[0336]  
[0337]       public func entity(  
[0338]            at point: CGPoint, in space: CoordinateSpace  
[0339]       ) -> Entity?  
[0340]       public func entities(  
[0341]            at point: CGPoint, in space: CoordinateSpace

```
[0342]    ) -> [Entity]

[0343]

[0344]    public func hitTest(

[0345]        point: CGPoint, in space: CoordinateSpace,

[0346]        query: CollisionCastQueryType, mask: CollisionGroup

[0347]    ) -> [CollisionCastHit]

[0348]    }

[0349]    '''

[0350]

[0351]    `

[0352]

[0353]    * The physics origin

[0354]    * The camera mode on a phone operating system (i.e., whether it uses AR)

[0355]    * The environment (e.g., a skybox)

[0356]    * The audio listener

[0357]    * Rendering and debug options

[0358]

[0359]

[0360]    @available(phoneOS, computerOS)

[0361]    extension Content {

[0362]        public typealias DebugOptions = ARView.DebugOptions
```

```

[0363]     public typealias RenderOptions = ARView.RenderOptions

[0364]     public typealias Environment = ARView.Environment

[0365]

[0366]     public var debugOptions: DebugOptions = []

[0367]     public var renderOptions: RenderOptions = []

[0368]     public var environment: Environment = .default

[0369]     public var physicsOrigin: Entity? = nil

[0370]     public var audioListener: Entity? = nil

[0371]

[0372]     public struct CameraMode {

[0373]         public static var nonAR: CameraMode

[0374]         public static var automaticAR: CameraMode

[0375]         public static func manualAR(_ sesion: ARSession) -> CameraMode

[0376]     }

[0377]

[0378]     public var cameraMode: CameraMode = .automaticAR

[0379] }

[0380] ```

```

[0381]

[0382] In some examples, the view is designed to provide processes with the option of *\*isolation\**: that is, whether its entities should operate in an independent simulation from any other entities in the application. For example, suppose a developer created two views in a scene of their app. If those views were both isolated, they would each run their own

independent set of systems, so behaviors like physics collisions would never occur between entities in different view hierarchies. If those views were *\*not\** isolated, however, they would both share a common set of systems, and thus behaviors like physics collisions could span entities across both views. In some examples, an isolated view would likely create its own Scene, while a non-isolated view would exist within a shared Scene.

[0383] In some examples, a flexible container view of the view with a frame can be explicitly set in the same way as any other view. In some examples, clients that wish to size a view to match the size of a containing entity can do so by manually calculating the bounding box of that entity and setting that as the view's frame, like so:

[0384]

[0385] struct CustomModel: View {

[0386] @State var modelSize: Size3D?

[0387]

[0388] var body: some View {

[0389] View { content in

[0390] if let model = try? await ModelEntity("airplane") {

[0391] content.add(model)

[0392] let bounds = model.visualBounds(relativeTo: nil)

[0393] modelSize = content.convert(bounds, to: .local).size

[0394] }

[0395] }.frame(width: modelSize?.width, height: modelSize?.height)

[0396] }

[0397] }

[0398] ```

```
[0399]
[0400]    /// A view for displaying 3D framework content.
[0401]    ///
[0402]    /// Use a view to display rich 3D content using the 3D framework in your app.
[0403]    /// A view passes a struct conforming to a protocol
[0404]    /// to the `make` and `update` closures, which you can use to add and remove
[0405]    /// 3D framework entities to your view.
[0406]    ///
[0407]    /// Here is a simple example showing how you can display a custom
[0408]    /// ``ModelEntity`` using view:
[0409]    ///
[0410]    ///     struct ModelExample: View {
[0411]    ///         var body: some View {
[0412]    ///             View { content in
[0413]    ///                 if let robot = try? await ModelEntity(named: "robot") {
[0414]    ///                     content.add(robot)
[0415]    ///                 }
[0416]    ///             Task {
[0417]    ///                 // perform any additional async work to configure
[0418]    ///                 // the content after the view is rendered
[0419]    ///             }
```

[0420]     ///  
[0421]     ///  
[0422]     ///  
[0423]     ///  
[0424]     /// Note that the closure used in the example above is `async`, and can be used  
[0425]     /// to load contents from your app's bundle or from any `URL` in the background.  
[0426]     /// While your content is loading, view will automatically display a  
[0427]     /// placeholder view, which you can customize using the optional `placeholder`  
[0428]     /// parameter. It is strongly recommended to load your content asynchronously to  
[0429]     /// avoid introducing a hang in your app.  
[0430]     ///  
[0431]     /// You can also use the optional `update` closure on your view to  
[0432]     /// update your 3D framework content in response to changes in your view's state.  
[0433]     ///  
[0434]     /// On a mixed reality operating system, view displays your 3D framework content  
inline in true 3D  
[0435]     /// space, occupying the available space in your app's 3D bounds. This is  
[0436]     /// represented by the content type. On a phone operating system,  
[0437]     /// view displays its content in an AR camera view by default, and can  
[0438]     /// display in a "non-AR" mode when requested or when AR or the device's  
camera  
[0439]     /// is unavailable. On a computer operating system, a view always displays its  
content in a

[0440]     /// non-AR mode. On both a phone operating system and a computer operating system, this is represented by the

[0441]     /// content type.

[0442]     ///

[0443]     /// View has a flexible size by default, and does not size itself based

[0444]     /// on the 3D framework content it is displaying.

[0445]     ///

[0446]     /// For more advanced uses of a 3D framework, such as subscribing to 3D framework

[0447]     /// events, performing coordinate conversions, or working with AR capabilities,

[0448]     /// refer to the Protocol types.

[0449]     @available(mixedrealityOS, phoneOS, computerOS)

[0450]     public struct View<Content: View>: View {}

[0451]

[0452]     @available(mixedrealityOS)

[0453]     extension View {

[0454]         /// Creates a View.

[0455]         ///

[0456]         /// - Parameters:

[0457]             ///     - `make`: A closure that you can use to set up and configure the

[0458]             ///     initial content in your view. This closure is `async` to

[0459]             ///     provide a natural place to load any initial state required to

[0460]       ///     display your view, like a `ModelEntity` loaded from disk.

[0461]       ///     It is strongly recommended to perform all loading operations

[0462]       ///     asynchronously to avoid creating a hang in your app's UI. While

[0463]       ///     your `make` closure is still evaluating, the `placeholder` view

[0464]       ///     will be displayed instead.

[0465]       ///     - `update`: An optional closure that you can use to update your

[0466]       ///     View's content in response to changes in your view's

[0467]       ///     state.

[0468]       ///     - `placeholder`: A view to display while your View's `make`

[0469]       ///     closure is being evaluated. For example, you could use a

[0470]       ///     `ProgressView` as a loading indicator.

[0471]     public init<P: View>(  
 [0472]         make: @Sendable @escaping @MainActor  
 [0473]         (inout Content) async -> Void,  
 [0474]         update: (@MainActor (inout Content) -> Void)? = nil,  
 [0475]         @ViewBuilder placeholder: () -> P  
 [0476]     ) where Content == Content.Body<P>  
 [0477]  
 [0478]     /// Creates a view.  
 [0479]     ///  
 [0480]     /// - Parameters:

[0481]       ///   - `make`: A closure that you can use to set up and configure the

[0482]       ///   initial content in your `view`. This closure is `async` to

[0483]       ///   provide a natural place to load any initial state required to

[0484]       ///   display your view, like a `ModelEntity` loaded from disk.

[0485]       ///   It is strongly recommended to perform all loading operations

[0486]       ///   asynchronously to avoid creating a hang in your app's UI. While

[0487]       ///   your `make` closure is still evaluating, a system-defined

[0488]       ///   placeholder view is displayed instead.

[0489]       ///   - `update`: An optional closure that you can use to update your

[0490]       ///   View's content in response to changes in your view's

[0491]       ///   state.

[0492]   public init(  
[0493]       make: @Sendable @escaping @MainActor  
[0494]       (inout Content) async -> Void,  
[0495]       update: (@MainActor (inout Content) -> Void)? = nil  
[0496]   ) where Content == Content.Body<DefaultPlaceholder>  
[0497]  
[0498]   public typealias DefaultPlaceholder = Placeholder  
[0499]   }  
[0500]  
[0501]   /// A view that represents the default placeholder for a view.

```
[0502]    ///
[0503]    /// You don't create this type directly. `View` creates values for you.
[0504]    @available(phoneOS
[0505]    @available(computerOS)
[0506]    @available(mixedrealityOS)
[0507]    public struct Placeholder: View {}
[0508]
[0509]    @available(phoneOS
[0510]    @available(computerOS)
[0511]    extension View {
[0512]        /// Creates a View.
[0513]        ///
[0514]        /// - Parameters:
[0515]        ///     - `make`: A closure that you can use to set up and configure the
[0516]        ///     initial content in your view. This closure is `async` to
[0517]        ///     provide a natural place to load any initial state required to
[0518]        ///     display your view, like a `ModelEntity` loaded from disk.
[0519]        ///     It is strongly recommended to perform all loading operations
[0520]        ///     asynchronously to avoid creating a hang in your app's UI. While
[0521]        ///     your `make` closure is still evaluating, the `placeholder` view
[0522]        ///     will be displayed instead.
```

```

[0523]    /// - `update`: An optional closure that you can use to update your
[0524]    ///    View's content in response to changes in your view's
[0525]    ///    state.
[0526]    /// - `placeholder`: A view to display while your `View`'s `make`
[0527]    ///    closure is being evaluated. For example, you could use a
[0528]    ///    `ProgressView` as a loading indicator.
[0529]    public init<P: View>(
[0530]        make: @Sendable @escaping @MainActor
[0531]        (inout Content) async -> Void,
[0532]        update: (@MainActor (inout Content) -> Void)? = nil,
[0533]        @ViewBuilder placeholder: () -> P
[0534]    ) where Content == Content.Body<P>
[0535]
[0536]    /// Creates a view.
[0537]    ///
[0538]    /// - Parameters:
[0539]    ///     - `make`: A closure that you can use to set up and configure the
[0540]    ///     initial content in your view. This closure is `async` to
[0541]    ///     provide a natural place to load any initial state required to
[0542]    ///     display your view, like a `ModelEntity` loaded from disk.
[0543]    ///     It is strongly recommended to perform all loading operations

```

```
[0544]    ///    asynchronously to avoid creating a hang in your app's UI. While
[0545]    ///    your `make` closure is still evaluating, a system-defined
[0546]    ///    placeholder view is displayed instead.
[0547]    ///    - `update`: An optional closure that you can use to update your
[0548]    ///    `View`'s content in response to changes in your view's
[0549]    ///    state.
[0550]    public init(
[0551]        make: @Sendable @escaping @MainActor
[0552]        (inout Content) async -> Void,
[0553]        update: (@MainActor (inout Content) -> Void)? = nil
[0554]    ) where Content == Content.Body<DefaultPlaceholder>
[0555]
[0556]    public typealias DefaultPlaceholder = Placeholder
[0557]    }
[0558]
[0559]    /// Represents a subscription to a 3D framework events, and provides the ability
[0560]    /// to cancel that subscription.
[0561]    @available(mixedrealityOS, phoneOS, computerOS)
[0562]    public struct EventSubscription {
[0563]
[0564]        /// Cancels the subscription to events.
```

```
[0565]    ///
[0566]    /// This method can be used to cancel the subscription before destruction.
[0567]    public func cancel()
[0568]    }
[0569]
[0570]    @available(mixedrealityOS, phoneOS, computerOS)
[0571]    /// An ordered, mutable collection of ``Entity``s
[0572]    public protocol EntityCollection : Collection where Element == Entity, Index ==
Int {
[0573]
[0574]    /// Adds the specified entity to the end of this collection.
[0575]    ///
[0576]    /// - Parameters:
[0577]    ///   - entity:      The entity to add to the collection.
[0578]    mutating func append(_ entity: Entity)
[0579]
[0580]    /// Adds the specified sequence of entities to the end of this collection,
[0581]    /// in order.
[0582]    ///
[0583]    /// - Parameters:
[0584]    ///   - sequence:    The entities to add to the collection.
```

[0585] mutating func append<S>(contentsOf sequence: S) where S: Sequence,  
S.Element: Entity

[0586]

[0587] /// Adds the specified entity to this collection directly before the entity

[0588] /// at the given index. If the entity is already located before the index,

[0589] /// the collection will not change.

[0590] ///

[0591] /// - Parameters:

[0592] /// - entity: The entity to add to the collection.

[0593] /// - index: The index of an entity to insert in front

[0594] /// of. If `endIndex` is provided, the entity

[0595] /// will be appended.

[0596] mutating func insert(\_ entity: Entity, beforeIndex index: Int)

[0597]

[0598] /// Adds the specified sequence of entities to this collection in order,

[0599] /// directly before the entity at the given index.

[0600] ///

[0601] /// - Parameters:

[0602] /// - sequence: A sequence of entities to add to the collection.

[0603] /// - index: The index of an entity to insert in front

[0604] /// of. If `endIndex` is provided, the

[0605] /// entities will be appended.

[0606] mutating func insert<S>(contentsOf sequence: S, beforeIndex index: Int) where  
S: Sequence, S.Element: Entity

[0607]

[0608] /// Removes the entity from the collection.

[0609] ///

[0610] /// - Parameters:

[0611] /// - entity: The entity to remove from the collection.

[0612] mutating func remove(\_ entity: Entity)

[0613]

[0614] /// Removes the entity at the given index from this collection.

[0615] ///

[0616] /// - Parameters:

[0617] /// - index: The index of the entity to remove from the collection.

[0618] mutating func remove(at index: Int)

[0619]

[0620] /// Removes all entities from this collection.

[0621] mutating func removeAll()

[0622]

[0623] /// Removes all entities from this collection that satisfy the given predicate.

[0624] mutating func removeAll(where: (Entity) throws -> Bool) rethrows

[0625]

[0626] /// Replaces all entities in this collection with those from the given

```

[0627]    /// sequence.

[0628]    ///

[0629]    /// - Parameters:

[0630]    /// - entities:          The sequence of entities that will replace

[0631]    ///                      the collection's current contents.

[0632]    mutating func replaceAll<S>(_ entities: S) where S: Sequence, S.Element:
Entity

[0633]    }

[0634]

[0635]    /// A protocol representing the content of a ``View``.

[0636]    ///

[0637]    /// You usually do not need to interface with this protocol directly. On
mixedrealityOS,

[0638]    /// you use ``ViewContent`` with your ``View``, while on phoneOS and

[0639]    /// computerOS, you use ``ViewCameraContent``.

[0640]    @available(mixedrealityOS, phoneOS, computerOS)

[0641]    public protocol ViewContentProtocol {

[0642]        /// The type of collection used for `entities`.

[0643]        associatedtype Entities: EntityCollection

[0644]

[0645]        /// The 3D framework entities to be displayed in this ``View``.

[0646]        var entities: Entities { get nonmutating set }

```

```

[0647]
[0648]    /// Subscribes to the provided 3D framework ``Event``.
[0649]    ///
[0650]    /// - Parameters:
[0651]    ///   - `event`: The 3D framework ``Event`` type to subscribe to.
[0652]    ///   - `sourceObject`: An optional ``EventSource`` to filter events to
[0653]    ///     (e.g., an entity). If no source is provided, the resulting
[0654]    ///     subscription will receive events for all objects in this
[0655]    ///     ``View``.
[0656]    ///   - `componentType`: An optional ``Component`` type to filter events
[0657]    ///     to. If no component type is provided, the resulting subscription
[0658]    ///     will receive events for all components in this ``View``, if
[0659]    ///     the event type is for a component.
[0660]    ///   - `handler`: A closure to run when the `event` occurs.
[0661]    ///
[0662]    /// - Returns: A value representing the subscription to this event stream.
[0663]    ///     Call ``cancel()`` on this value to end the subscription.
[0664]    func subscribe<E: Event>(to event: E.Type,
[0665]                            on sourceObject: EventSource?,
[0666]                            componentType: Component.Type?,
[0667]                            _ handler: @escaping (E) -> Void) -> EventSubscription

```

```
[0668]    }

[0669]

[0670]    @available(mixedrealityOS, phoneOS, computerOS

[0671]    extension ViewContentProtocol {

[0672]        /// Adds `entity` to this content.

[0673]        public func add(_ entity: Entity)

[0674]

[0675]        /// Removes `entity` from this content, if present.

[0676]        public func remove(_ entity: Entity)

[0677]

[0678]        /// Subscribes to the provided 3D framework ``Event``.

[0679]        ///

[0680]        /// - Parameters:

[0681]        ///   - `event`: The 3D framework ``Event`` type to subscribe to.

[0682]        ///   - `sourceObject`: An optional ``EventSource`` to filter events to

[0683]        ///     (e.g., an entity). If no source is provided, the resulting

[0684]        ///     subscription will receive events for all objects in this

[0685]        ///     ``View``.

[0686]        ///   - `handler`: A closure to run when the `event` occurs.

[0687]        ///

[0688]        /// - Returns: A value representing the subscription to this event stream.
```

```

[0689]      ///      Call ``cancel()`` on this value to end the subscription.

[0690]      public func subscribe<E: Event>(to event: E.Type,

[0691]          on sourceObject: EventSource?,

[0692]          _ handler: @escaping (E) -> Void) -> EventSubscription

[0693]

[0694]      /// Subscribes to the provided 3D framework ``Event``.

[0695]      ///

[0696]      /// - Parameters:

[0697]      ///   - `event`: The 3D framework ``Event`` type to subscribe to.

[0698]      ///   - `componentType`: An optional ``Component`` type to filter events

[0699]      ///     to. If no component type is provided, the resulting subscription

[0700]      ///     will receive events for all components in this ``View``, if

[0701]      ///     the event type is for a component.

[0702]      ///   - `handler`: A closure to run when the `event` occurs.

[0703]      ///

[0704]      /// - Returns: A value representing the subscription to this event stream.

[0705]      ///      Call ``cancel()`` on this value to end the subscription.

[0706]      public func subscribe<E: Event>(to event: E.Type,

[0707]          componentType: Component.Type? = nil,

[0708]          _ handler: @escaping (E) -> Void) -> EventSubscription

[0709]      }

```

[0710]

[0711]     /// A collection of entities in a ``View``.

[0712]     ///

[0713]     /// This collection is used by ``ViewContentProtocol/entities``.

[0714]     @available(mixedrealityOS, phoneOS, computerOS)

[0715]     public struct ViewEntityCollection: EntityCollection

[0716]     {

[0717]         // Note: boilerplate details of EntityCollection protocol conformance

[0718]         // omitted for brevity

[0719]     }

[0720]

[0721]     /// The content of a ``View`` that is displayed inline.

[0722]     ///

[0723]     /// On mixedrealityOS, `ViewContent` is used to display your 3D framework  
content

[0724]     /// inline in true 3D space, occupying the available space in your app's 3D

[0725]     /// bounds.

[0726]     ///

[0727]     /// You can use `ViewContent` to add and remove entities, subscribe to

[0728]     /// 3D framework events, and perform coordinate conversions between 3D  
framework

[0729]     /// entity space and a 2D framework View's coordinate space.

```
[0730]    @available(mixedrealityOS)

[0731]    public struct ViewContent: ViewContentProtocol {

[0732]        public var entities: ViewEntityCollection { get nonmutating set }

[0733]

[0734]        public func subscribe<E: Event>(to event: E.Type,

[0735]                                       on sourceObject: EventSource?,

[0736]                                       componentType: Component.Type?,

[0737]                                       _ handler: @escaping (E) -> Void) -> EventSubscription

[0738]

[0739]        /// The default view contents of a ``View`` using

[0740]        /// ``ViewContent``.

[0741]        ///

[0742]        /// You don't create this type directly. ``View`` creates values for

[0743]        /// you.

[0744]        public struct Body<Placeholder: View>: View {}

[0745]    }

[0746]

[0747]    @available(mixedrealityOS)

[0748]    extension ViewContent {

[0749]        /// Convert the provided point from an entity to a 2D framework coordinate

[0750]        /// space.
```

[0751]       /// - Parameters:

[0752]       ///   - `point`: The point to be converted. This point is interpreted in

[0753]       ///       3D framework entity coordinates. This means the value is expected to

[0754]       ///       be expressed in meters, with the Y axis pointed up, and expressed

[0755]       ///       relative to the entity's origin.

[0756]       ///   - `entity`: The entity that `point` should be interpreted in. If no

[0757]       ///       entity is provided, the value is interpreted to be in the

[0758]       ///       coordinate space of the ``ViewContent``.

[0759]       ///   - `space`: The 2D framework coordinate space in which the resulting point

[0760]       ///       should be expressed.

[0761]       ///

[0762]       /// - Returns: A converted point expressed in 2D framework coordinates. This

[0763]       ///       means the value will be expressed in points, with the Y axis pointing

[0764]       ///       down, and expressed relative to the coordinate space's

[0765]       ///       top-leading-back origin.

[0766]       public func convert(

[0767]            \_point: SIMD3<Float>, from entity: Entity? = nil, to space:

CoordinateSpace

[0768]       ) -> Point3D

[0769]

[0770]       /// Convert the provided rect from an entity to a SwiftUI coordinate space.

[0771]       ///

[0772]       /// - Parameters:

[0773]       ///    - `bounds`: The bounds to be converted. These bounds are interpreted in

[0774]       ///       3D framework entity coordinates. This means the value is expected to

[0775]       ///       be expressed in meters, with the Y axis pointed up.

[0776]       ///    - `entity`: The entity that `rect` should be interpreted in. If no

[0777]       ///       entity is provided, the value is interpreted to be in the

[0778]       ///       coordinate space of the ``ViewContent``.

[0779]       ///    - `space`: The 2D framework coordinate space in which the resulting rect

[0780]       ///       should be expressed.

[0781]       ///

[0782]       /// - Returns: A converted rect expressed in 2D framework coordinates. This means

[0783]       ///    the value will be expressed in points, with the Y axis pointing down,

[0784]       ///    and expressed relative to the coordinate space's top-leading-back

[0785]       ///    origin.

[0786]       public func convert(

[0787]            \_bounds: BoundingBox, from entity: Entity? = nil, to space: CoordinateSpace

[0788]       ) -> Rect3D

[0789]

[0790]       /// Convert the provided point from a 2D framework coordinate space to an entity.

[0791]       ///

[0792]       /// - Parameters:

[0793]       ///   - `point`: The point to be converted. This point is interpreted in

[0794]       ///       2D framework coordinates. This means the value is expected to be

[0795]       ///       expressed in points, with the Y axis pointing down, and expressed

[0796]       ///       relative to the coordinate space's top-leading-back origin.

[0797]       ///   - `space`: The 2D framework coordinate space in which `point` should be

[0798]       ///       interpreted.

[0799]       ///   - `entity`: The entity that the resulting point should be expressed

[0800]       ///       in. If no entity is provided, the value is interpreted to be in

[0801]       ///       the coordinate space of the ``ViewContent``.

[0802]       ///

[0803]       /// - Returns: A converted point expressed in 3D framework entity coordinates.

[0804]       ///   This means the value will be expressed in meters, with the Y axis

[0805]       ///       pointed up, and expressed relative to the entity's origin.

[0806]       public func convert(  
[0807]            \_point: Point3D, from space: CoordinateSpace, to entity: Entity? = nil  
[0808]       ) -> SIMD3<Float>  
[0809]

[0810]       /// Convert the provided rect from a SwiftUI coordinate space to an entity.

[0811]       ///

[0812]       /// - Parameters:

[0813]       ///   - `rect`: The rect to be converted. This rect is interpreted in

[0814]       ///       2D framework coordinates. This means the value is expected to be

[0815]       ///       expressed in points, with the Y axis pointing down, and expressed

[0816]       ///       relative to the coordinate space's top-leading-back origin.

[0817]       ///   - `space`: The 2D framework coordinate space in which `rect` should be

[0818]       ///       interpreted.

[0819]       ///   - `entity`: The entity that the resulting rect should be expressed

[0820]       ///       in. If no entity is provided, the value is interpreted to be in

[0821]       ///       the coordinate space of the ``ViewContent``.

[0822]       ///

[0823]       /// - Returns: A converted bounding box expressed in 3D framework entity

coordinates.

[0824]       ///   This means the value will be expressed in meters, with the Y axis

[0825]       ///   pointed up, and expressed relative to the entity's origin.

[0826]       public func convert(

[0827]           \_rect: Rect3D, from space: CoordinateSpace, to entity: Entity? = nil

[0828]       ) -> BoundingBox

[0829]       }

[0830]

[0831]       /// The content of a ``View`` that is displayed through a camera.

[0832]       ///

[0833]       /// On phoneOS, `ViewCameraContent` displays content in an AR camera view by

[0834]     /// default, and can display in a "non-AR" mode when requested or when AR or  
the

[0835]     /// device's camera is unavailable. On computerOS, `ViewCameraContent` always

[0836]     /// displays its content in a non-AR mode.

[0837]     ///

[0838]     /// You can use `ViewCameraContent` to add and remove entities, subscribe

[0839]     /// to 3D framework events, configure the AR environment, and perform  
coordinate

[0840]     /// conversions such as projections and raycasts between the ``View``'s

[0841]     /// space and a 2D framework View coordinate space.

[0842]     @available(phoneOS)

[0843]     @available(computerOS)

[0844]     public struct ViewCameraContent: ViewContentProtocol {

[0845]

[0846]         public typealias DebugOptions = ARView.DebugOptions

[0847]         public typealias RenderOptions = ARView.RenderOptions

[0848]         public typealias Environment = ARView.Environment

[0849]

[0850]         /// Options for drawing overlay content in a scene that can aid in

[0851]         /// debugging.

[0852]         public var debugOptions: DebugOptions

[0853]

[0854]       /// Options for enabling / disabling specific rendering effects.

[0855]       public var renderOptions: RenderOptions

[0856]

[0857]       /// Description of background, lighting and acoustic properties for your AR

[0858]       /// scene.

[0859]       public var environment: Environment

[0860]

[0861]       /// Selects the entity which defines the origin of the physics simulation in

[0862]       /// the scene.

[0863]       ///

[0864]       /// By default, the origin of the physics simulation coincides with the

[0865]       /// origin of the scene. However, in many AR applications, the physics

[0866]       /// simulation should be relative to a certain anchor in the scene.

[0867]       ///

[0868]       /// For example: An AR app simulates a game of Jenga. The Jenga blocks are

[0869]       /// added as entities with rigid body components to a scene. The app detects

[0870]       /// an anchor (marker or plane) at which the game should be placed in

[0871]       /// real world. The physics simulation should be relative to the detected

[0872]       /// anchor. The anchor can be moved around in the real world without

[0873]       /// affecting the physics simulation.

[0874]       ///

[0875]       /// To achieve this:

[0876]       /// 1. Add a new entity to the scene that tracks the anchor position.

[0877]       /// 2. Set `physicsOrigin` to the entity to indicate that this entity's

[0878]       ///   transform determines the origin of the physics simulation.

[0879]       /// 3. Optionally, parent the Jenga blocks to the anchor entity. This way

[0880]       ///   the Jenga blocks are automatically updated when the anchor position

[0881]       ///   changes.

[0882]       ///

[0883]       /// Example:

[0884]       /// View { content in

[0885]       ///   // Define your anchor entity and add it to the scene.

[0886]       ///   let myAnchor = AnchorEntity(.image(group: "References", name:

[0886]       ///   "GameImage"))

[0887]       ///   content.add(myAnchor)

[0888]       ///

[0889]       ///   // Set myAnchor as the origin of the physics simulation.

[0890]       ///   content.physicsOrigin = myAnchor

[0891]       ///

[0892]       ///   // Add the simulated blocks to the anchor.

[0893]       ///   myAnchor.children.append(block0)

[0894]       ///   myAnchor.children.append(block1)

[0895]       /// }

[0896]       /// ``

[0897]       /// Using this setup, all forces, velocities, etc. will be simulated

[0898]       /// relative to `myAnchor`. Moving the anchor will not affect the

[0899]       /// simulation.

[0900]       public var physicsOrigin: Entity?

[0901]

[0902]       /// The entity which defines the listener position and orientation for

[0903]       /// spatial audio.

[0904]       ///

[0905]       /// By default this `audioListener` property is nil, which means that the

[0906]       /// active camera entity will be used as the audio listener. The

[0907]       /// `audioListener` can be set to any entity in the `scene` to use the

[0908]       /// transform of the entity as the audio listener position and orientation.

[0909]       public var audioListener: Entity?

[0910]

[0911]       public enum CameraMode {

[0912]           /// Use non-AR rendering, with a background defined by the `environment`

[0913]           /// property.

[0914]       case nonAR

[0915]

[0916]       /// Use a default ``ARSession`` that is created and automatically

```
[0917]    /// managed by the view after initialization. This is the default value.
[0918]    case automaticAR
[0919]
[0920]    /// Provide a custom ``ARSession`` that powers this ``View``'s AR
[0921]    /// rendering.
[0922]    ///
[0923]    /// If you have an existing or custom session, set it as the view's
[0924]    /// session to replace the default session and run it by calling
[0925]    /// `session.run(_:options:)`.
[0926]    case manualAR(ARSession)
[0927]    }
[0928]
[0929]    public var cameraMode: CameraMode
[0930]
[0931]    /// The default view contents of a ``View`` using ``ViewCameraContent``.
[0932]    ///
[0933]    /// You don't create this type directly. ``View`` creates values for you.
[0934]    public struct Body<Placeholder: View>: View {}
[0935]    }
[0936]
[0937]    @available(phoneOS)
```

[0938]     @available(computerOS)

[0939]     extension ViewCameraContent {

[0940]         /// Projects a point `point` from the 3D world coordinate system of the

[0941]         /// scene to the 2D pixel coordinate system of the viewport.

[0942]         public func project(  
[0943]             vector: SIMD3<Float>, to space: CoordinateSpace  
[0944]         ) -> CGPoint?

[0945]

[0946]         /// Unproject a 2D point from the view onto a plane in 3D world coordinates.

[0947]         ///

[0948]         /// A 2D point in the view's coordinate space can refer to any point along a

[0949]         /// line segment in the 3D coordinate space. Unprojecting gets the 3D

[0950]         /// position of the point along this line segment that intersects the

[0951]         /// provided plane.

[0952]         ///

[0953]         /// - Parameters:

[0954]         ///     - `point`: A point in the provided coordinate space.

[0955]         ///     - `space`: The coordinate space in which `point` should be

[0956]         ///     interpreted.

[0957]         ///     - `planeTransform`: The transform used to define the coordinate

[0958]         ///     system of the plane. The coordinate system's positive Y axis is

[0959]       ///     assumed to be the normal of the plane.

[0960]       ///     - `relativeToCamera` : If the plane transform is relative to camera

[0961]       ///     space or world space. Defaults to false.

[0962]       ///

[0963]       /// - Returns: 3D position in world coordinates or nil if unprojection is

[0964]       ///     not possible.

[0965]       public func unproject(  
[0966]           point: CGPoint, from space: CoordinateSpace, ontoPlane: float4x4,  
[0967]           relativeToCamera: Bool = false  
[0968]       ) -> SIMD3<Float>?

[0969]

[0970]       /// Unprojects a 2D point `point` in the pixel coordinate system of

[0971]       /// `viewport`, converting it to a view-space 3D coordinate.

[0972]       public func unproject(  
[0973]           point: CGPoint, viewport: CGRect  
[0974]       ) -> SIMD3<Float>?

[0975]

[0976]       /// Calculates a ray in the AR scene that corresponds to point on screen.

[0977]       ///

[0978]       /// - Parameters:

[0979]       ///     - `point`: A point in the provided coordinate space.

[0980]       ///   - `space`: The coordinate space in which `point` should be

[0981]       ///       interpreted.

[0982]       ///

[0983]       /// - Returns: A ray in AR scene coordinates that goes from the camera

[0984]       ///       origin through the specified `point` on screen.

[0985]       public func ray(  
[0986]           through: CGPoint, in space: CoordinateSpace  
[0987]       ) -> (origin: SIMD3<Float>, direction: SIMD3<Float>)?

[0988]

[0989]       /// Gets the closest entity in the AR scene at the specified point on

[0990]       /// screen.

[0991]       ///

[0992]       /// - Parameters:

[0993]       ///   - `point`: A point in the provided coordinate space.

[0994]       ///   - `space`: The coordinate space in which `point` should be

[0995]       ///       interpreted.

[0996]       ///

[0997]       /// - Returns: The entity at `point`. Returns `nil` if no entity was found.

[0998]       ///

[0999]       /// - Important: Hit tests (ray-casts) are performed against the collision

[1000]       ///   shapes. Entities without a proper `CollisionComponent` (see

[1001]       ///    `HasCollision`) are ignored in hit tests!

[1002]       public func entity(  
[1003]           at point: CGPoint, in space: CoordinateSpace  
[1004]       ) -> Entity?  
[1005]  
[1006]       /// Gets all entities in the AR scene at the specified point on screen.  
[1007]       ///  
[1008]       /// - Parameters:  
[1009]       ///    - `point`: A point in the provided coordinate space.  
[1010]       ///    - `space`: The coordinate space in which `point` should be  
[1011]       ///       interpreted.  
[1012]       ///  
[1013]       /// - Returns: A list of entities at `point`. Returns an empty array if no  
[1014]       ///       entities were found.  
[1015]       ///  
[1016]       /// - Important: Hit tests (ray-casts) are performed against the collision  
[1017]       ///       shapes. Entities without a proper `CollisionComponent` (see  
[1018]       ///       `HasCollision`) are ignored in hit tests!  
[1019]       public func entities(  
[1020]           at point: CGPoint, in space: CoordinateSpace  
[1021]       ) -> [Entity]

[1022]

[1023]     /// Searches the scene for entities at the specified point in the view.

[1024]     ///

[1025]     /// - Parameters:

[1026]     ///     - `point`: A point in the provided coordinate space.

[1027]     ///     - `space`: The coordinate space in which `point` should be

[1028]     ///     interpreted.

[1029]     ///     - `query`: The query type.

[1030]     ///     - `mask`: The collision mask. This value can be used to prevent

[1031]     ///     hits with certain objects. The default value is

[1032]     ///     `CollisionGroup.all`, which means the ray can hit all objects. See

[1033]     ///     `CollisionFilter` for details.

[1034]     ///

[1035]     /// - Returns: An array of hit-test results.

[1036]     ///

[1037]     /// - Important: Hit-tests are performed against the collision shapes.

[1038]     ///     Entities without a proper `CollisionComponent` (see `HasCollision`)

[1039]     ///     are ignored in hit tests!

[1040]     ///

[1041]     /// To make a hit-test against real-world objects call `hitTest(\_types:)`.

[1042]     public func hitTest(

```
[1043]         point: CGPoint, in space: CoordinateSpace,
[1044]         query: CollisionCastQueryType, mask: CollisionGroup
[1045]     ) -> [CollisionCastHit]
[1046] }
[1047] ```
[1048]     Here is a simple example showing how physics can be configured within a
`View`, observing collision events as they occur and updating the UI.
[1049]
[1050]     struct PhysicsExample: View {
[1051]         @State var collisionCount = 0
[1052]         @State var subscription: EventSubscription?
[1053]
[1054]         var body: some View {
[1055]             View { content in
[1056]                 let floor = ModelEntity(mesh: .generateBox(width: 1, height: 1, depth: 1))
[1057]                 floor.collision = CollisionComponent(shapes: [.generateBox(width: 1,
height: 1, depth: 1)])
[1058]                 floor.physicsBody = PhysicsBodyComponent(shapes: [.generateBox(size:
[1,1,1]), mass: 1, material: nil, mode: .static)
[1059]                 floor.position.y = -0.6
[1060]                 content.add(floor)
[1061]
[1062]                 let box = ModelEntity(mesh: .generateBox(size: 0.1))
```

```

[1063]         box.collision = CollisionComponent(shapes: [.generateBox(size:
[0.1,0.1,0.1]))])

[1064]         box.physicsBody = PhysicsBodyComponent(shapes: [.generateBox(size:
[0.1,0.1,0.1])), mass: 1)

[1065]         box.position.y = 0.3

[1066]         content.add(box)

[1067]

[1068]         subscription = content.subscribe(to: CollisionEvents.Began.self) {

[1069]             collisionCount += 1

[1070]         }

[1071]     }

[1072]     .overlay(alignment: .bottom) {

[1073]         Text("Collision Count: \$(collisionCount)")

[1074]         .padding()

[1075]         .background(.regularMaterial, in: Capsule())

[1076]     }

[1077] }

[1078] }

[1079] ```

```

[1080] Here is a complete example of an app loading a 3D model asynchronously, showing a progress indicator while waiting and an error message if the model could not be loaded.

```

[1081] struct AsyncModelExample: View {

```

```
[1082]      @State var error: Error?

[1083]

[1084]      var body: some View {

[1085]          if let error {

[1086]              Text("Error! \ \(error.localizedDescription)")

[1087]          } else {

[1088]              View { content in

[1089]                  do {

[1090]                      let model = try await ModelEntity(named: "computer system
_closed")

[1091]

[1092]                      model.transform.rotation = simd_quatf(angle: .pi/4, axis: [0,1,0])

[1093]                      model.transform.translation.y = -0.08

[1094]

[1095]                      let lid = model.findEntity(named: "lid") as? ModelEntity

[1096]                      lid?.transform.rotation = simd_quatf(angle: -1.919, axis: [1,0,0])

[1097]                  } catch {

[1098]                      self.error = error

[1099]                  }

[1100]              } placeholder: {

[1101]                  ProgressView()

[1102]              }
```

[1103] }

[1104] }

[1105] }

[1106] ```

[1107] An alternative to `View` as proposed is to instead model it as a protocol. One possible version of a protocol-based approach is described as `ViewRepresentable`. In some examples, there is value in having *\*both\** an inline closure-based API as proposed herein and a protocol-based API.

[1108] The simple use cases of `View` are easier and more convenient to express as a view than as a protocol. For example:

[1109]

[1110] View { content in

[1111] content.add(ModelEntity(...))

[1112] }

[1113] ```

[1114]

[1115] Or, if we added a convenience initializer that takes an entity directly, as noted in the Future Directions section:

[1116]

[1117] ```swift

[1118] View(modelEntity)

[1119] ```

[1120]

[1121] As opposed to a protocol-based approach, which requires extracting the view out into a separate type:

[1122] ```swift

[1123] struct ModelView: ViewProtocol {

[1124] func makeContent(\_ content: inout Content) {

[1125] content.add(ModelEntity(...))

[1126] }

[1127] }

[1128] ```

[1129] In some examples, it is important to be able to apply modifiers directly to a `View` without having to wrap it in its own distinct type. With the other Representable types, a 2D framework and an application framework provide enough functionality that it's reasonable to expect these views to typically be "islands" that don't require a ton of back-and-forth with another 2D framework modifiers or views. In some examples with a 3D framework, that isn't the case. In some examples, 2D framework-based UI content can be mixed in a 3D framework hierarchy, using the `View` symbols API proposed. This API uses an additional `ViewBuilder` parameter that naturally extends the `View` closure-based syntax, similar to the `Canvas` API:

[1130]

[1131] View { content in

[1132] ...

[1133] } symbols: {

[1134] ...

[1135] }

[1136] ```

[1137]

[1138] Other potential extensions noted in the \*Future Directions\* section are worth

[1139] considering too, like control over the `View`'s scope, e.g.:

[1140]

[1141] `View(scope: .isolated) { content in ... }`

[1142] `````

[1143]

[1144] Or supporting different rendering styles like portals, e.g.:

[1145]

[1146] `View(style: .portal, cameraPosition: Point3D(...)) { content in ... }`

[1147] `````

[1148] These sorts of extensions are also representable using a protocol, but can be more elegantly and concisely expressed using a single `View` type with a variety of initializers.

[1149] In a similar vein to the protocol-based `View` described above, 2D framework for mixedrealityOS has for some time supported 3D framework integration via an `EntityRepresentable` protocol modeled after an existing 2D framework and an application framework Representable types (e.g., `UIViewRepresentable`). While there is value in treating 3D framework entities as a peer to `UIView`'s and `NSView`'s, we have in practice found this model to be a poor fit for 3D framework for a few reasons:

[1150] Clunky layout semantics:\*\* Unlike 2D framework and application framework, 3D framework has no notion of a layout system, and thus there is no good intrinsic size to use for an `EntityRepresentable`. This could be partially addressed by exposing a public `sizeThatFits` API to allow developers to provide a fitting size, defaulting to a flexible size. This would still result in sizing semantics that differ from the other Representable types, however. 3D framework lacks UI frameworks capabilities:\*\* Unlike the other Representable types, 3D framework entities do not have a gesture system, a responder chain for event

handling, a notion of environment / preferences for hierarchical data flow, and other features common to 2D framework , an application framework, and another 3D framework . We have retrofitted some of these capabilities for entities in 2D framework , but we have no intention of shipping this as API. \*This is not a bad thing —\*3D framework does not need these tools for the use cases it is built for — but it implies a different design for how to integrate 3D framework content in a 2D framework hierarchy, as well as how to mix 2D framework content in a 3D framework hierarchy (the subject of a separate proposal). Poor cross-platform compatibility:\*\* An `EntityRepresentable` API is unlikely to be a good fit to bring to phoneOS and computerOS. While 3D framework is a more foundational framework for content on mixedrealityOS, the same is not true of phoneOS and computerOS. An API structure like `EntityRepresentable` implies a level of deep integration with 2D framework that is not present on those platforms. Further, many phoneOS and computerOS APIs may not make sense in an `EntityRepresentable` context, and the semantics of an `EntityRepresentable` API seem like a poor fit for a top-level entry point into 3D framework , where APIs like anchoring are expected to be used instead.

[1151]     ### Split View into multiple views

[1152]     Instead of consolidating around a single `View` type, we could publish multiple new Views that offer different APIs and functionality based on their use cases. For example, we could add an `EntityView` type designed for the “shared scope” scenario, displaying an entity hierarchy inline:

[1153]     EntityView { entity in

[1154]         entity.addChild(ModelEntity(...))

[1155]     }

[1156]     ``

[1157]

[1158]     The entity provided to clients would be an Entity subclass with additional

[1159]     “top-level” functionality like the ability to subscribe to events, convert

[1160]     coordinate spaces to and from 2D framework , etc.

[1161]

[1162] We would also add a `View` variant that is instead tailored to the

[1163] “isolated scope” scenario, providing a `3D framework.Scene` directly:

[1164]

[1165] ```swift

[1166] View { scene in

[1167] scene.addEntity(ModelEntity(...))

[1168] }

[1169] ```

[1170] Similarly, this would also provide a `3D framework.Scene` subclass with additional functionality. Additional View types could be added in the future for other use cases, like a `RealityCamera` type for displaying a camera into another scene.

[1171] This approach would allow us to reuse existing 3D framework abstractions. However, it has a few downsides: It further bifurcates the API between mixedrealityOS and phoneOS / computerOS. We are likely not able to support isolation or multiple scenes within an app for v1, nor are we likely to support shared entity hierarchies for some time on phoneOS or computerOS. Therefore, `View` as described above would only be available on phoneOS and computerOS, while `EntityView` would only be available on mixedrealityOS. This would hopefully be lifted eventually, but would take some time. There is likely not enough API differentiation between these cases. Clients of `EntityView` are likely to need many of the same APIs as `View`, like coordinate conversions, event subscriptions, and the ability to mix in UI with the entity hierarchy (to be addressed in a follow-up proposal). It is difficult to name these APIs clearly, in a way that justifies the differentiation between them. Developers new to 3D framework may find having both an `EntityView` and a `View` to be confusing, especially since the primary decision factor on when to use one or the other is based on whether the developer wants a shared or isolated scope, something that is not communicated in the names themselves. As the number of use cases grows over time, this

approach may require the introduction of many new views, e.g. a ``Portal`` type for hosting 3D framework content in a portal.

[1172] Previous versions of this proposal have advocated for either a 3D framework Scene or Entity at the root of a ``View`` instead of the proposed ``ViewContentProtocol`` protocol. There are a few different forms this can take that we'll enumerate:

[1173] Using a 3D framework Scene (specifically a subclass) uniformly within a ``View`` would lean into the Scene being a “namespace” for a ``View``, and would effectively always enforce an isolated scope for a ``View``. This was rejected for a few reasons: It is likely not doable in the mixedrealityOS 1.0 timeframe, as it requires support for multiple scenes in 3D framework. Scenes in 3D framework are somewhat heavyweight, so this would not work well in cases where performance is important, e.g. if many ``View``s were displayed in a grid in an app's UI. The ``3D framework.Scene`` type conflicts with the ``2D framework.Scene`` type, and thus requires an explicit module qualifier whenever it is referenced in any code that imports both frameworks. This makes ``3D framework.Scene`` an awkward type to use within 2D framework, and thus we should avoid leaning too heavily into abstractions that use it.

[1174] Using an Entity subclass uniformly within a ``View`` would leave open the possibility that ``View`` could use either a shared or isolated scope, while also reusing an existing 3D framework abstraction. This seems more tractable than using Scene uniformly, but was also rejected: It requires introducing several “top-level” concepts on this Entity subclass that were previously on Scene, like the ability to subscribe to events, and particularly on phoneOS and computerOS, ray-casting into the scene, controlling its ARSession, etc. This blurs the lines between an Entity and a Scene in 3D framework, which could be confusing to developers. Some APIs are quite awkward to introduce directly on an Entity type. To name two examples: The follow-up proposal to mix UI content in a ``View`` may propose using “symbols” to represent UI that is provided directly from a ``View`` `ViewBuilder`, which can then be added to the developer's entity hierarchy. This API would be strange to use from an Entity subclass, as clients would access symbols from an ``entity.symbols`` property and then need to add each symbol to the hierarchy — even though the API would suggest that the symbols are already a part of the root entity. In the future we may wish to allow apps to provide an existing 3D framework Scene directly to a ``View``. This may compose awkwardly with a root Entity type provided by the ``View``. It enforces a single uniform interface across all ``View``s. This limits future API flexibility to restrict functionality based on context — for

example, if some APIs should only be available when in an isolated scope, like controlling the physics origin.

[1175] One other option could be to provide either an Entity or a Scene to the body of a `View` depending on some context or signal. For example, based on an explicit scope:

[1176]

[1177] `View(scope: .shared) { entity in ... }`

[1178] `View(scope: .isolated) { scene in ... }`

[1179] `...`

[1180] This would improve API flexibility and allow a more appropriate type to be provided based on context. However, this approach suffers from many of the same downsides as described in the *\*Split `View` into multiple views\** alternative above. Further, the problems described there seem worse in this case, as there is a surprising API jump occurring within the same View type. This jump is far less noticeable with the proposed `ViewContentProtocol` approach, as the terminology used is consistent, with common functionality captured directly in the protocol, and in most cases occurs transparently to the developer.

[1181] In addition, this differentiation would be quite awkward in v1, where mixedrealityOS would only support a shared scope and phoneOS / computerOS would only support an isolated scope. This means the default initializers for `View` would use different types on each platform, further bifurcating code across platforms. Some other names were considered in the proposed API.

[1182] Alternatives include: `ViewState`, `ViewContext`, `ViewSubscene`

[1183] Note that nesting types within the `View` type was rejected due to the use of a generic type, making it impossible for clients to write types like `View.Content`, as they would need to provide a concrete type to `View` like `View<ViewContent>.Content`.

[1184] The names for the various `ViewContent` types would ideally be less verbose, with names like `Content` and `CameraContent`. These names seemed a bit too general to

warrant dropping the namespace prefix, however, and the hope is that clients will rarely if ever have to use these types directly.

[1185] Alternatives to ``ViewCameraContent`` were considered, like ``ViewProjectedContent`` and ``ViewportContent``, which also express the general concept of a 2D projection into a 3D environment, both in AR and non-AR contexts. Including “camera” in the name seemed like the most straightforward way to encapsulate this, and also offers consistent terminology with some of its APIs like ``cameraMode``.

[1186] Other names that lean more on the “AR” capabilities were also considered, like ``ViewARCameraContent`` and ``ViewARContent``. These were rejected to allow the same API to express a non-AR render-to-texture mode on mixedrealityOS in the future.

[1187] Alternatives to ``ViewContent`` were also considered, like ``ViewInlineContent`` (as proposed in a previous version of this proposal), ``ViewEmbeddedContent``, or simply ``RealityContent``. We ultimately landed on ``ViewContent`` to convey the primacy of using 3D framework content embedded in a true 3D context.

[1188] ``ViewContentProtocol`` was named ``ViewContent`` in a previous version of this proposal, but we ultimately determined that using the ``Protocol`` suffix would be the most appropriate way to preserve the primacy of the concrete ``ViewContent`` type while still conveying the notion that all conformances are kinds of content for a ``View``. The ``Protocol`` suffix is somewhat unfortunate, but this does follow a precedent of other protocols in Swift like ``StringProtocol``, as well as in frameworks like 2D framework with ``CoordinateSpaceProtocol``.

[1189] A previous version of this proposal included support for 3D framework raycasting via ``ViewCameraContent`` on phoneOS, providing parity with ``View`` on phoneOS. We have opted to remove this support from the proposal for now while we take more time to determine the right shape of that API and align better with future directions for 3D framework on phoneOS.

[1190] A strong goal of this 3D framework and 2D framework integration effort is to enable “80% 3D framework , 20% 2D framework ” apps: that is, apps that are primarily simulation-driven, but would like to mix in some UI. To support these cases, we need to provide apps with the tools to host 2D framework Views within a ``View``’s entity hierarchy.

Extending the notion of "symbols" as used in 2D framework `Canvas` API for inserting 2D framework Views as 3D framework entities.

[1191] While `View` is intended to compose well with the existing 2D framework Gesture APIs, there are a few pieces that require new APIs: The ability to determine which entity was hit-tested for a given gesture, Conveniences to convert a given gesture's value into 3D framework coordinates, A lower-level gesture primitive to receive raw touches for use in some clients that need more direct access to this data, like third-party game engines. These will be addressed in follow-up proposals.

[1192] The proposed `View` API is great for lightweight uses of 3D framework , and while it can scale up to complex use cases, we believe there is merit to also providing a more advanced API that scales better to these use cases. In particular, a protocol-based API could help: Improve ergonomics of asynchronously-loaded state, avoiding the optional dance required when using async state stored as 2D framework `@State` loaded in a `.task` modifier, Providing better facilities for handling errors, e.g. should a loading operation throw an error, Serving as a natural scope for more complex imperative code using 3D framework .

[1193] We are investigating a supplemental protocol for addressing these advanced uses of `View`, tentatively termed `ViewRepresentable`, which could look something like this:

[1194] @MainActor

[1195] protocol ViewRepresentable: View where Body == Never {

[1196] associatedtype AsyncState = Void

[1197] associatedtype Content: ViewContentProtocol = ViewDefaultContent

[1198] func makeState() async throws -> AsyncState

[1199] func makeContent(\_ content: inout Content, state: inout AsyncState)

[1200] func updateContent(\_ content: inout Content, state: inout AsyncState)

[1201] }

[1202] #if os(mixedrealityOS)

[1203]     typealias ViewDefaultContent = ViewContent

[1204]     #else

[1205]     typealias ViewDefaultContent = ViewCameraContent

[1206]     #endif

[1207]     ``

[1208]     This protocol is analogous to the proposed closure-based structure of ``View``, with a few extra features. Most notably, this includes support for loading async state required before the view can be shown via the ``makeState()`` function and ``AsyncState`` associated type. This state is then passed in to the ``makeContent`` and ``updateContent`` functions, in a similar fashion to the ``Cache`` in the ``Layout`` protocol. ``makeState()`` is also marked ``throws``, allowing clients to propagate errors up to the outer view if desired. A ``ViewRepresentable`` could then be used either as a ``View`` directly, translating it into a ``View`` behind the scenes, or by using some special ``callAsFunction``-type view builders to support use of a placeholder or [symbols](<https://github.pie.apple.com/TDG/SwiftUI-Evolution/pull/12>).

[1209]     For example, a simple ``ViewRepresentable`` could be defined as:

[1210]     ```swift

[1211]     struct ExampleRepresentable: ViewRepresentable {

[1212]         @Observable class AsyncState {

[1213]             var model: ModelEntity

[1214]         }

[1215]         func makeState() async throws -> AsyncState {

[1216]             let model = try await ModelEntity("...")

[1217]             return AsyncState(model: model)

[1218]         }

[1219]       func makeContent(\_ content: inout Content, state: inout AsyncState) {

[1220]           content.add(state.model)

[1221]       }

[1222]    }

[1223]    ``

[1224]

[1225]    And used either directly as a view:

[1226]

[1227]    ```swift

[1228]       struct ExampleView: View {

[1229]           var body: some View {

[1230]               ExampleRepresentable()

[1231]       }

[1232]    }

[1233]    ``

[1234]    Or using ViewBuilder-like syntax, where the representable is given access to a View proxy view to which modifiers like ``gesture`` can be applied:

[1235]    ```swift

[1236]       struct ExampleView: View {

[1237]           var body: some View {

[1238]               ExampleRepresentable { state, view in

[1239]                   view.gesture(...)

[1240]        }  
[1241]        }  
[1242]        }  
[1243]        ```\n\n[1244]        As well as the ability to provide a placeholder view and / or symbols:\n\n[1245]        ```swift\n\n[1246]        struct ExampleView: View {\n\n[1247]            var body: some View {\n\n[1248]                ExampleRepresentable { state, view in\n\n[1249]                    view.gesture(...)\n\n[1250]                } placeholder: {\n\n[1251]                    ProgressView()\n\n[1252]                } symbols: { state in\n\n[1253]                    @Bindable \$state = state\n\n[1254]                    Toggle("Enabled", isOn: \$state.foo)\n\n[1255]                    .tag("toggle")\n\n[1256]                }\n\n[1257]            }\n\n[1258]        }\n\n[1259]        ```\n

[1260] A `ViewRepresentable` could be initialized with additional parameters as needed, and could work with dynamic properties like `@State` and `@Binding` directly. This allows additional state to be passed in from the containing view; for example:

```
[1261] struct ExampleRepresentable: ViewRepresentable {  
[1262]     // mutable state that the representable could write to, e.g.  
[1263]     // if a collision event occurs  
[1264]     @Binding var showAlien: Bool  
[1265]     // immutable state passed from the view  
[1266]     var showDebugUI: Bool  
[1267]     ...  
[1268] }  
[1269] struct ExampleView: View {  
[1270]     @State var showAlien = false  
[1271]     @State var showDebugUI = false  
[1272]     var body: some View {  
[1273]         ExampleRepresentable(  
[1274]             showAlien: $showAlien, showDebugUI: showDebugUI  
[1275]         ) { state, view in  
[1276]             view.gesture(...)   
[1277]         }  
[1278]         Toggle("Show Debug UI", isOn: $showDebugUI)  
[1279]     }
```

[1280] }

[1281] ```

[1282] Error handling could also be provided by passing a `Phase` into the view closure, similar to the `AsyncImage` API:

[1283] struct ExampleView: View {

[1284] @State var showAlien = false

[1285] var body: some View {

[1286] ExampleRepresentable(showAlien: \$showAlien) { phase, view in

[1287] switch phase {

[1288] case .success(let state):

[1289] view.gesture(...)

[1290] case .failure(let error):

[1291] Text(error.localizedDescription)

[1292] case .empty:

[1293] ProgressView()

[1294] }

[1295] } symbols: { state in

[1296] Toggle("Show Alien", isOn: \$showAlien)

[1297] .toggleStyle(.button)

[1298] }

[1299] }

[1300] }

[1301]     '''

[1302]     This direction is still under consideration and may be proposed as a follow-up to the core `View` API.

[1303]     As noted earlier, we would like to give apps control over a `View`'s isolation: that is, whether it maintains its own independent set of systems that are unaware of other entities in the app. Once 3D framework can support the ability to host multiple Scenes within an app, we would like to expose control over this in `View`. One way this could be done is by adding new initializers to optionally provide a scope, like so:

[1304]     View(scope: .isolated) { content in

[1305]         ...

[1306]     }

[1307]     '''

[1308]     We would likely want to enforce the use of a static scope, as dynamically changing the scope is unlikely to be possible or desirable. We could do so through the use of generics, creating new scope types like so:

[1309]     public protocol ViewScope { ... }

[1310]     public struct IsolatedViewScope: ViewScope {}

[1311]     public struct SharedViewScope: ViewScope {}

[1312]

[1313]     extension ViewScope where Self == IsolatedViewScope {

[1314]         public static var isolated: IsolatedViewScope { .init() }

[1315]     }

[1316]     extension ViewScope where Self == SharedViewScope {

[1317]         public static var shared: SharedViewScope { .init() }

[1318] }

[1319] ```

[1320]

[1321] And extending `View` to make use of them:

[1322]

[1323] extension View {

[1324]     public typealias Scope = ViewScope

[1325]

[1326]     public init<S: Scope>(

[1327]         scope: S,

[1328]         make: @escaping (inout Content) -> (),

[1329]         update: ((inout Content) -> ())? = nil

[1330]     ) where Content == ViewScopedContent<S> {

[1331]         content = Content(make: make, update: update)

[1332]     }

[1333] }

[1334] ```

[1335] Another form of lighter-weight isolation would be to allow custom Systems to query for entities by `View`, such that they can group certain behaviors per `View`. This could potentially be done by providing an opaque namespace ID to `ViewContent` like so:

[1336] extension ViewContent {

[1337]     public var namespace: AnyHashable { get }

```
[1338]    }
[1339]    ```
[1340]
[1341]    Which could then be queried on the Scene:
[1342]
[1343]    class MySystem: System {
[1344]        func update(context: SceneUpdateContext) {
[1345]            // For a specific View:
[1346]            var targetNamespace = ...
[1347]            context.scene.performQuery(
[1348]                EntityQuery(where: .namespace(matches: targetNamespace)
[1349]            ) { entity in
[1350]                // ...
[1351]            }
[1352]
[1353]            // For each View:
[1354]            context.scene.performQuery(
[1355]                GroupedEntityQuery(by: .namespace)
[1356]            ) { entities in
[1357]                // ...
[1358]            }
```

[1359] }

[1360] }

[1361] ``

[1362]

[1363] Further exploration is needed here. For now, developers can work around this without any new API by relying on a custom component:

[1364]

[1365] struct MyRootComponent: TransientComponent, Codable {}

[1366]

[1367] struct MyView: View {

[1368] var body: some View {

[1369] View { content in

[1370] var rootEntity = Entity()

[1371] rootEntity.components.set(MyRootComponent())

[1372] // Add mores entities:

[1373] rootEntity.children.append(...)

[1374] content.add(rootEntity)

[1375] }

[1376] }

[1377] }

[1378]

[1379] class MySystem: System {

```

[1380]     func update(context: SceneUpdateContext) {
[1381]         // For each View:
[1382]         context.scene.performQuery(EntityQuery(where:
.has(MyRootComponent.self))) { rootEntity in
[1383]             // For each entity in View:
[1384]             context.scene.performQuery(EntityQuery(where: .isDescendant(of:
rootEntity))) { entity in
[1385]                 ...update View entities...
[1386]             }
[1387]         }
[1388]     }
[1389] }
[1390] ```

```

[1391] Some apps will want to render their 3D content within a 2D surface, using either a perspective-aware portal: `![Portal](portal-example.png)` Or a render-to-texture technique: `![Render-to-Texture](render-to-texture-example.png)`

[1392] These use cases require the use of an isolated 3D framework Scene under the hood. We could potentially represent this using a new ``style`` property on ``View``, like so:

```

[1393] View(style: .portal) { content in ... }
[1394] View(style: .viewport) { content in ... }
[1395] ```

```

[1396] The style type would be statically typed and used to infer available properties of the content. To control the perspective from which the content is viewed, an optional ``cameraPosition`` property could be provided:

[1397]     View(style: .portal, cameraPosition: Point3D(...)) { content in ... }

[1398]     '''

[1399]     The layout semantics of a `View` in this configuration would be slightly different, as it would occupy zero depth. The `viewport` style for render-to-texture use cases would reuse the `ViewCameraContent` type from phoneOS and computerOS, while the `portal` style may be best expressed with a separate `ViewContentProtocol` type, or could reuse the same content type as well.

[1400]     Some apps may want the ability to display a camera with another view into the same shared scene. For example, consider this mockup for a car configurator app: ![Car Configurator App Mockup](car-configurator-app-mockup.png)

[1401]     Here cameras into the shared car scene are used in the configurator's UI. This could be done by, for example, providing a `PerspectiveCamera` to a `View` like so:

[1402]     View(viewing: perspectiveCamera)

[1403]     '''

[1404]

[1405]     Or perhaps by constructing the `View` with an existing Scene type:

[1406]

[1407]     View(viewing: sharedScene) { content in

[1408]         content.camera = PerspectiveCamera()

[1409]     }

[1410]     '''

[1411]     Some apps simply need to display a single entity inline in their hierarchy. For this case, we could provide a simple variant of `View` that directly takes an entity:

[1412]     '''swift

[1413] View(entity)

[1414] ``

[1415] The primary concern with this approach is that it could imply to developers that the `View` will be sized automatically to match the size of the entity, which would not be the case and is unlikely to be feasible for reasons described earlier. However, the convenience of this API could still be worthwhile, provided these sizing semantics are communicated clearly.

[1416] An often-requested enhancement to 3D framework is to provide a declarative, value-type interface for constructing entity hierarchies, similar to 2D framework views. This would fit in nicely with `View`, enabling developers to use a consistent declarative programming style throughout their hierarchy. Such an API is quite an ambitious undertaking, however, and requires solving a number of challenges, like: Reconciling existing reference types (like `Entity`) with new declarative versions (e.g., what names can or should they use?), Supporting mechanisms like `@State` in a way that works with 3D framework, Interoperation with other ECS-based concepts like custom Systems

[1417]

[1418] One approach that was attempted was the use of Swift Concurrency and the

[1419] AsyncSequence to utilise 3D framework Events. This was the proposed API

[1420] solution for that:

[1421] public protocol ViewContentProtocol {

[1422] func events<E: Event>(<E>)

[1423] for event: E.Type, on sourceObject: EventSource?,

[1424] componentType: Component.Type?

[1425] ) -> EventSequence<E>

[1426] }

[1427] extension ViewContentProtocol {

```

[1428]     public func events<E: Event>(
[1429]         for event: E.Type, on sourceObject: EventSource? = nil
[1430]     ) -> EventSequence<E>
[1431]     public func events<E: Event>(
[1432]         for event: E.Type, componentType: Component.Type?
[1433]     ) -> EventSequence<E>
[1434] }
[1435] public struct EventSequence<E: Event>: AsyncSequence {}
[1436] ``

```

[1437] When attempting this implementation, it was discovered that 3D framework event handlers run into issues when scheduled by the Swift async runtime. For example, we could not represent a 3D framework event subscription as an ``AsyncSequence``. This is because when an event is fired 3D framework guarantees that it will suspend execution until the event has been received by all subscriptions. When you ``await`` on a value from an ``AsyncSequence``, you may receive the value after an indeterminate amount of time (as decided by the 2D framework scheduler). So there would be no way to implement this guarantee. For example, using the current proposed solution:

```

[1438]
[1439] struct RocketComponent: Component {
[1440]     private var isLaunched: Bool = false
[1441]     private var isInOrbit: Bool = false
[1442]
[1443]     public func launch() {
[1444]         isLaunched = true

```

```
[1445]         print("Rocket has launched!")
[1446]     }
[1447]
[1448]     public func moveToOrbit() {
[1449]         precondition(isLaunched) // Will not crash in this example
[1450]         isInOrbit = true
[1451]         print("Rocket is now orbiting Earth")
[1452]     }
[1453] }
[1454] struct RocketshipSimulator: View {
[1455]     @State var state = SimulationState()
[1456]     @State var subscription: EventSubscription?
[1457]     var body: some View {
[1458]         View { content in
[1459]             let rocket = await state.loadRocket()
[1460]             content.entities += rocket
[1461]
[1462]             // Every time a RocketComponent is added, launch the rocket.
[1463]             subscription = content.subscribe(to: ComponentEvents.DidAdd.self,
[1464]                                             componentType: RocketComponent.self) { event in
[1465]                 event.entity.components[RocketComponent.self]!.launch()
```

[1466] }

[1467]

[1468] // launch() will be called via the subscription before set() returns

[1469] rocket.components.set(RocketComponent())

[1470]

[1471] // moveToOrbit's precondition will pass because launch() was called

[1472] rocket.components[RocketComponent.self]!.moveToOrbit()

[1473] }

[1474] }

[1475] }

[1476] ```

[1477] In some examples, techniques describe a 2D framework for hosting custom content within apps on operating system. It is designed to scale to support a variety of use cases for working with content. One key use case that we aim to enable on operating systems is the ability to embed UI within a scene. This is crucial for turning a passive 3D scene into an interactive experience, and for building true volumetric apps on operating systems. A 2D framework is not expected to provide all the tools needed to build rich 3D interfaces and content on its own in operating system, and thus a tight integration story between UI built in 2D framework and content authored in 3D framework will be essential to allow developers to explore this space and build a new generation of 3D applications.

[1478] In some examples, 3D apps should be able to: Make use of the full suite of UI components, including both informational UI like text and info panels, and interactive controls like buttons and sliders, Place their UI content anywhere in their 3D scene (e.g., pinning a UI panel relative to the selected piece of a 3D model), Automatically get the system-standard look-and-feel for their controls, Support the standard hand-based interaction model on operating system, including indirect touch (gaze-and-pinch), direct touch, and their accessibility variants.

[1479] In some examples, a new concept of "attachments" to 3D framework for hosting UI content within its entity hierarchy. Attachments would only be available on operating system at first. Here is a simple example of how attachments are used:

```
[1480] struct MyCoolSimulation: View {  
[1481]     @State var state = SimulationState()  
[1482]  
[1483]     enum AttachmentID: Int {  
[1484]         case button  
[1485]     }  
[1486]  
[1487]     var body: some View {  
[1488]         View { content, attachments in  
[1489]             content.add(state.modelEntity)  
[1490]             if let button = attachments.entity(for: AttachmentID.button) {  
[1491]                 button.position = SIMD3(x: 0.25, y: 0.1, z: 0.1)  
[1492]                 content.add(button)  
[1493]             }  
[1494]         } attachments: {  
[1495]             Button("Bigger!", action: {  
[1496]                 state.modelEntity.scale *= 2  
[1497]             })  
[1498]         }.tag(AttachmentID.button)
```

[1499] }

[1500] }

[1501] }

[1502]

[1503] @Observable

[1504] class SimulationState {

[1505] var modelEntity = ModelEntity(mesh: .generateSphere(radius: 0.1))

[1506] }

[1507] ``

[1508] Attachments are declared using the optional ``attachments`` `ViewBuilder`. Each attachment should be given a tag to identify it using the ``tag()`` modifier. Within the view's body, the new ``attachments`` parameter provides access to attachment entities via the ``entity(for:)`` function, keyed by the attachment view's tag. Because these attachments are vended as entities, the developer can treat them as they would any other content. Note that these entities are *\*not\** automatically added to the entity hierarchy; developers are expected

[1509] to manage this themselves so that they are consciously placed within the ``view``.

[1510] Attachments can be dynamically added and removed from the ``attachments`` `ViewBuilder`, including support for 2D framework transition and animation system. Under the hood, attachments are rendered the same way other 2D framework views are rendered, only with a 3D position based on how the developer places the attachment entities.

[1511] attachment entities can also be inspected more directly through the use of the

[1512] new ``ViewAttachmentComponent``:

[1513]

[1514] public struct ViewAttachmentComponent: TransientComponent {

[1515]       /// The bounding box of the attachment, expressed in meters.

[1516]       public var bounds: BoundingBox { get }

[1517]

[1518]       /// The identifier used for this attachment.

[1519]       public var id: AnyHashable { get }

[1520]       }

[1521]       '''

[1522]

[1523]       In some examples, the only properties exposed on `ViewAttachmentComponent` are: 1. Its bounds, allowing apps to place a attachment relatively based on its size. 2. Its ID, allowing custom Systems to associate the attachment entity with corresponding UI state.

[1524]       Attachment entities are instances of the new `ViewAttachmentEntity` class, which provides a more direct, type-safe accessor to the underlying `ViewAttachmentComponent`:

[1525]       public class ViewAttachmentEntity: Entity {

[1526]       public var attachment: ViewAttachmentComponent { get }

[1527]       }

[1528]       '''

[1529]       Attachment entities do not expose most of their implementation details, like their underlying meshes or materials. They do not participate in physics collisions by default, though apps can opt them in to physics by manually adding a collider (whose shape must also be calculated manually).

[1530]       @available(mixedrealityOS)

[1531]       extension View {

[1532]       /// Creates a `View`.

[1533]       ///

[1534]       /// - Parameters:

[1535]       ///    - `make`: A closure that you can use to set up and configure the

[1536]       ///       initial content in your `View`.

[1537]       ///    - `update`: An optional closure that you can use to update your

[1538]       ///       `View`'s content in response to changes in your view's

[1539]       ///       state.

[1540]       ///    - `attachments`: A `ViewBuilder` of attachment views to use in your

[1541]       ///       `View`'s content. Access entities representing your

[1542]       ///       attachments via the ``ViewAttachments`` passed in

[1543]       ///       to your `make` and `update` closures using the

[1544]       ///       ``ViewAttachments/entity(for:)`` function.

[1545]       ///

[1546]       /// - Note: Attachment views are not automatically added to your

[1547]       ///    `View`'s content. Instead, you must explicitly add these

[1548]       ///       entities to the hierarchy. Attachment entities can be placed normally

[1549]       ///       or parented to any other entity. Access details of a attachment

[1550]       ///       entity, like its bounds, using methods on the ``ViewAttachmentEntity``

[1551]       ///       type.

[1552]       public init<A: View>(  
 [1553]           make: @Sendable @escaping @MainActor

[1554] (inout RealityViewContent, ViewAttachments)

[1555] async -> Void,

[1556] update: (@MainActor (inout ViewContent, ViewAttachments) -> Void)? =  
nil,

[1557] @ViewBuilder attachments: @escaping () -> A

[1558] ) where Content == AttachmentContent<A, DefaultPlaceholder>

[1559]

[1560] /// Creates a `View`. ///

[1561] /// - Parameters:

[1562] /// - `make`: A closure that you can use to set up and configure the

[1563] /// initial content in your `View`. This closure is `async` to

[1564] /// provide a natural place to load any initial state required to

[1565] /// display your `View`, like a `ModelEntity` loaded from disk.

[1566] /// It is strongly recommended to perform all loading operations

[1567] /// asynchronously to avoid creating a hang in your app's UI. While

[1568] /// your `make` closure is still evaluating, the `placeholder` view

[1569] /// will be displayed instead.

[1570] /// - `update`: An optional closure that you can use to update your

[1571] /// `View`'s content in response to changes in your view's

[1572] /// state.

[1573] /// - `attachments`: A `Builder` of attachment views to use in your

[1574] /// `View`'s content. Access entities representing your

```

[1575]    /// attachments via the ``ViewAttachments`` passed in
[1576]    /// to your `make` and `update` closures using the
[1577]    /// ``ViewAttachments/entity(for:)`` function.
[1578]    /// - `placeholder`: A view to display while your `View`'s `make`
[1579]    /// closure is being evaluated. For example, you could use a
[1580]    /// `ProgressView` as a loading indicator.
[1581]    ///
[1582]    /// - Note: Attachment views are not automatically added to your
[1583]    /// `View`'s content. Instead, you must explicitly add these
[1584]    /// entities to the hierarchy. Attachment entities can be placed normally
[1585]    /// or parented to any other entity. Access details of a attachment
[1586]    /// entity, like its bounds, using methods on the ``ViewAttachmentEntity``
[1587]    /// type.
[1588]    public init<A: View, P: View>(
[1589]        make: @Sendable @escaping @MainActor
[1590]        (inout ViewContent, ViewAttachments)
[1591]        async -> Void,
[1592]        update: (@MainActor (inout ViewContent, ViewAttachments) -> Void)? =
nil,
[1593]        @ViewBuilder placeholder: () -> P,
[1594]        @ViewBuilder attachments: @escaping () -> A

```

```

[1595]     ) where Content == AttachmentContent<A, P> public typealias
AttachmentContent<A, P> =

[1596]         ViewAttachmentContent<A, ViewContent.Body<P>>

[1597]     }

[1598]

[1599]     /// A view that represents the contents of a ``View`` with attachments.

[1600]     ///

[1601]     /// You don't create this type directly. ``View`` creates values for you.

[1602]     @available(mixedrealityOS)

[1603]     public struct ViewAttachmentContent<Attachments: View, Content: View>:
View {}

[1604]

[1605]     @available(mixedrealityOS)

[1606]     public struct ViewAttachments {

[1607]         /// Gets the identified attachment view as an entity, if the view with that

[1608]         /// tag exists.

[1609]         ///

[1610]         /// Note: resolving a attachment entity does not automatically add it to

[1611]         /// your ``View``'s content. You must add it yourself using a

[1612]         /// function like ``ViewContent/add(_):``.

[1613]         ///

[1614]         /// - Parameter id: The value that you used to tag the view when you

```

```

[1615]    /// define it in the `attachments` parameter of the ``View``
[1616]    /// initializer ``View/init(make:update:attachments:)``.
[1617]    /// - Returns: The resolved attachment entity, or `nil` if ``View``
[1618]    /// can't find an attachment view with the given `id`.
[1619]    public func entity(for id: some Hashable) -> ViewAttachmentEntity?
[1620]    }
[1621]    /// A component containing additional information about a view attachment
[1622]    /// entity provided via the ``ViewAttachments/entity(for:)``
[1623]    /// function.
[1624]    @available(mixedrealityOS 1.0)
[1625]    public struct ViewAttachmentComponent: TransientComponent, Identifiable {
[1626]        /// The bounding box of the view attachment, expressed in meters.
[1627]        public var bounds: BoundingBox { get }
[1628]
[1629]        /// The identifier used for this view attachment.
[1630]        public var id: AnyHashable { get }
[1631]    }
[1632]
[1633]    /// An entity that has a view attachment.
[1634]    @available(OS
[1635]    public class ViewAttachmentEntity: Entity {

```

[1636]       /// The view attachment component for this entity.

[1637]       public var attachment: ViewAttachmentComponent { get }

[1638]       }

[1639]       ```\n

[1640]       Below are a couple of complete examples on how attachments can be used. First,  
a

[1641]       view that asynchronously loads a robot model and adds a button anchored to its  
left hand:

[1642]

[1643]       struct InteractiveRobotExample: View {

[1644]        @State var state = SimulationState()

[1645]

[1646]        enum AttachmentID: Int {

[1647]         case button

[1648]       }

[1649]

[1650]       var body: some View {

[1651]         RealityView { content, attachments in

[1652]         await state.load()

[1653]         guard let robot = state.robot else { return }

[1654]         content.add(robot)

[1655]

```
[1656]         if let button = attachments.entity(for: AttachmentID.button),
[1657]             let hand = robot.findEntity(named: "left_hand")
[1658]         {
[1659]             // place the button to the left of the hand
[1660]             button.position.x = button.bounds.max.x - 0.05
[1661]             hand.addChild(button)
[1662]         }
[1663]     } placeholder: {
[1664]         ProgressView()
[1665]     } attachments: {
[1666]         Button("Wave!", action: state.wave)
[1667]         .tag(AttachmentID.button)
[1668]     }
[1669] }

[1670]
[1671] @Observable
[1672] class SimulationState {
[1673]     var robot: ModelEntity?
[1674]
[1675]     func load() async {
[1676]         guard robot == nil else { return }
```

```
[1677]         robot = try? await ModelEntity(named: "robot")
[1678]     }
[1679] }
[1680]
[1681] func wave() {
[1682]     if let robot, let animation = robot.availableAnimations
[1683]         .first(where: { $0.name == "wave" })
[1684]     {
[1685]         robot.playAnimation(animation)
[1686]     }
[1687] }
[1688] }
[1689] }
[1690] ```
[1691] Second, a view displaying a dynamic set of shapes, with UI to add new shapes and
[1692] delete each shape:
[1693]
[1694] struct InteractiveShapesExample: View {
[1695]     @State var state = SimulationState()
[1696]
[1697]     var body: some View {
```

```
[1698]      VStack {
[1699]          GeometryReader3D { geometry in
[1700]              simulationContent(in: geometry)
[1701]          }
[1702]          Button("Add Shape", action: state.addShape)
[1703]      }
[1704]  }

[1705]  func simulationContent(in geometry: GeometryProxy3D) -> some View {
[1706]      RealityView { _ in
[1707]          // empty to start
[1708]      } update: { content, attachments in
[1709]          // Add any new shapes to the hierarchy
[1710]          for (id, shape) in state.shapes where shape.parent == nil {
[1711]              content.add(shape)
[1712]
[1713]              let position = Point3D(
[1714]                  x: .random(in: 0...1), y: .random(in: 0...1),
[1715]                  z: .random(in: 0...1)
[1716]              ) * Point3D(geometry.size)
[1717]              shape.position = content.convert(position, from: .local)
[1718]
```

```
[1719]         if let delete = attachments.entity(for: id) {
[1720]             shape.addChild(delete)
[1721]         }
[1722]     }
[1723] } attachments: {
[1724]     ForEach(state.shapeIDs, id: \.self) { id in
[1725]         Button(action: { state.removeShape(id) }) {
[1726]             Image(systemName: "xmark")
[1727]         }.tag(id)
[1728]     }
[1729] }
[1730] }
[1731]
[1732] @Observable
[1733] class SimulationState {
[1734]     private var counter = 0
[1735]     var shapes = [Int: ModelEntity]()
[1736]
[1737]     var shapeIDs: [Int] { shapes.keys.sorted() }
[1738]
[1739]     func addShape() {
```

```
[1740]         let shape = ModelEntity(  
[1741]             mesh: .generateSphere(radius: Float.random(in: 0.1...0.4)),  
[1742]             materials: [UnlitMaterial(color: randomColor())])  
[1743]         shape.generateCollisionShapes(recursive: true) // make it hit-testable  
[1744]         shapes[counter] = shape  
[1745]         counter += 1  
[1746]     }  
[1747]  
[1748]     func removeShape(_ id: Int) {  
[1749]         if let shape = shapes[id] {  
[1750]             shape.removeFromParent()  
[1751]             shapes[id] = nil  
[1752]         }  
[1753]     }  
[1754]  
[1755]     private func randomColor() -> UIColor {  
[1756]         .init(  
[1757]             hue: .random(in: 0...1),  
[1758]             saturation: .random(in: 0.5...1),  
[1759]             brightness: 1, alpha: 1)  
[1760]     }
```

[1761] }

[1762] }

[1763] ``

[1764] Attachments are designed to be high-level and abstract away many of its implementation details. It is thus flexible and can be implemented in one of several ways, and can change over time as needed. However, to fully understand this proposal it can be useful to review the current direction for its implementation on Operating System.

[1765] When the developer provides a attachment view to `View`, a *\*proxy\** entity will be created for that view. That entity is what is provided to the `View`'s body via the `attachments.entity(for:)` function. The developer can parent that entity within their hierarchy and position it as they see fit. However, the underlying 2D framework content for this attachment is not directly parented to this proxy entity. Instead, `View` maintains a shallow hierarchy of attachment views, parenting them directly to the `View`. It then tracks the transforms of the corresponding proxy entities to keep in sync with the actual positions of those attachment views. To maintain good performance, updates to attachment view positions are coalesced into a single transaction.

[1766] However, there are some implications to consider: \* Apps cannot directly introspect the entity or layer hierarchy of the attachment entity. We view this as a benefit, as it reduces developer temptation to try and mess with that hierarchy in ways that will break over time. Attachment entities do not participate in physics collisions by default, and apps must manually calculate their colliders using the `ViewAttachmentComponent`'s bounds. (The latter could also be solved via smarter integration with `Entity.generateCollisionShapes()`) Certain hierarchical effects in 3D framework, like 3D clipping, would require an alternate means to propagate correctly, likely using a server-side mechanism. Such mechanisms already exist internally, but would need to be generalized to use with attachment entities. Note that there are currently no hierarchical effects as public API in 3D framework (and only a few as SPI), but this may one day change.

[1767] </details>

[1768]

[1769] A component-based approach is the leading alternative candidate, as it fits the most neatly into existing design patterns. This would work similarly to the

[1770] API for UICollectionView interop, and could look like so:

[1771]

[1772] struct MyCoolSimulation: View {

[1773] @StateObject var state = SimulationState()

[1774]

[1775] var body: some View {

[1776] View { content in

[1777] content.add(state.modelEntity)

[1778]

[1779] let button = Entity()

[1780] button[UIHostingComponent.self] = UIHostingComponent {

[1781] Button("Bigger!", action: {

[1782] state.modelEntity.scale \*= 2

[1783] })

[1784] }

[1785] button.position = SIMD3(x: 0.25, y: 0.1, z: 0.1)

[1786] content.add(button)

[1787] }

[1788] }

[1789] }

[1790]

[1791]     class SimulationState: ObservableObject {

[1792]         var modelEntity = ModelEntity(mesh: .generateSphere(radius: 0.1))

[1793]     }

[1794]     ``

[1795]     While this approach is more convenient in some ways, it presents some significant challenges: \* It is \*much\* harder to implement. It requires that we support the use of a fully heterogeneous and nesting 2D framework-3D framework hierarchy. This means that all UI framework features, like propagating the environment and preferences, coordinate conversions, gestures and hit testing, and so on must learn to work with this mixed entity-layer hierarchy. Doing so breaks many assumptions held in frameworks throughout the stack, and adds significant complexity to 2D frameworks and 3D frameworks to implement properly. We have attempted to do this before, and learned the hard way just how much of an added burden this places on our frameworks. Teaching 2D framework to interoperate smoothly in a heterogeneous hierarchy is already a massive implementation burden on 2D framework; while supporting this was a necessary and worthwhile tradeoff for 2D framework to take on at the time, we have more freedom with 3D framework to choose a different and simpler path. The aforementioned implementation burden is needed to support things like the environment propagating to entities within a `View` that use a `UIHostingComponent`, or gestures within that `UIHostingComponent` composing correctly with gestures from above in the hierarchy. These API promises must be upheld for `UIHostingComponent` to support the full feature set of a 2D framework, and thus the implementation burden must be maintained long-term as well.

[1796]     2D framework and 3D framework go hand in hand on operating system. The system is built from the ground up with 3D framework as a foundation by leveraging advanced 3D rendering techniques, deep integration with animations, and powerful simulation systems. Similarly, input on the operating system relies on 3D framework for 3D ray casting and collision detection. However, the bulk of this work has focused on extending existing gestures into 3D.

[1797] Because of this work, OS supports all the standard gestures vended by 2D frameworks. Techniques herein offer more expressive 3D support through extensions to existing gesture values like a `location3D` property on `DragGesture`. The introduction of `View` offers a powerful jumping off point from 2D framework into 3D framework. 2D framework integration can now host an arbitrary hierarchy, and we imagine developers will use this API to add a wide range of 3D functionality into their applications, from small flourishes of 3D in a primarily windowed app, all the way to fully immersive VR applications built almost entirely within 3D framework.

[1798] 3D framework's prominence on OS demands additional support from the gesture system. As applications lean more heavily on `View`, they begin to lose the ability to build interactions using 2D framework gesture system since `View` is somewhat opaque from 2D framework perspective. Similar to `Canvas`, it's not the case that every targetable item inside of 3D framework is mapped 1-1 with a `View`, which is where 2D gestures must be attached. Because of this, developers are incentivized to drop down to virtual kits for low level hand data to build their own custom interaction and hit testing systems. And because we are not exposing gaze data to applications, it is simply not possible to implement a wide range of interactions that are supported in 2D frameworks because of their deeper integration with the render server.

[1799] To bring the power of 2D gestures to `View` content a gesture modifier named `targetedToEntity` that allows the developer to associate a gesture with the primary entity targeted by the `Gesture`. The gesture's values, including [the new extensions to existing gestures] are modified to append the targeted entity and adds the ability to convert the gesture's spatial data between 2D and 3D coordinate spaces by leveraging the new `RealityCoordinateSpaceConverting` protocol. The gesture is also failed if it targets non-3D framework content, or if it targets the `_wrong_` 3D framework content, which we'll see in some of the examples below.

[1800] In this example, a sphere within a `View` is able to receive input and activate a `SpatialTapGesture` when tapped. The tap location is converted into the coordinate space of the sphere and used to place a new child sphere at that location.

[1801]

[1802] struct SpatialTapSphere: View {

```
[1803]     var body: some View {  
[1804]         View { content in  
[1805]             let sphere = ModelEntity(  
[1806]                 mesh: .generateSphere(radius: 0.1),  
[1807]                 materials: [SimpleMaterial(color: .blue, roughness: 0.25, isMetallic:  
false)],  
[1808]                 collisionShape: .generateSphere(radius: 0.1),  
[1809]                 mass: 1  
[1810]             )  
[1811]             sphere.physicsBody!.mode = .static  
[1812]             sphere.components.set(InputTargetComponent())  
[1813]             content.add(sphere)  
[1814]         }  
[1815]         .gesture(SpatialTapGesture().targetedToAnyEntity().onEnded { value in  
[1816]  
[1817]             // Convert the drag location from 2D local space of the 3D  
[1818]             // to the targeted entity's local space.  
[1819]             let tapPoint = value.convert(  
[1820]                 value.location3D,  
[1821]                 from: .local,  
[1822]                 to: value.entity  
[1823]             )
```

```

[1824]
[1825]         // Create a sphere at the tap location and add it as a child of the
[1826]         // targeted entity.
[1827]         let newEntity = ModelEntity(mesh: .generateSphere(radius: 0.01),
[1828]         materials: [SimpleMaterial(color: .red, roughness: 0.25, isMetallic:
false)])
[1829]         newEntity.position = tapPoint
[1830]         entity.addChild(newEntity)
[1831]     })
[1832] }
[1833] }
[1834] ```
[1835]

```

[1836] In the above example, `targetedToAnyEntity()` means that the modified gesture should only activates if it is operating on any `Entity`. We also propose allowing the developer to specify a particular `Entity` to restrict the modified gesture. For example, a developer could add a `MagnifyGesture` to the whole `View` while having a `DragGesture` only trigger when operating on a specific entity or any entity downstream from it using `targetedToEntity(entity)`:

```

[1837]
[1838] ```swift
[1839] struct ArcadeCabinet: View {
[1840]     @GestureState var dragLocation: Point3D = .zero
[1841]     @GestureState var inProgressScaleFactor: CGFloat = 1

```

```
[1842]      @State var cabinetScale: CGFloat = 1
[1843]      let joystick: Entity = makeJoystickEntity()
[1844]      var body: some View {
[1845]          RealityView { content in
[1846]              let screen = ScreenEntity()
[1847]              let cabinet = CabinetEntity()
[1848]              let redButton = ArcadeButtonEntity(.red)
[1849]              let greenButton = ArcadeButtonEntity(.green)
[1850]              joystick.setParent(cabinet)
[1851]              redButton.setParent(cabinet)
[1852]              greenButton.setParent(cabinet)
[1853]              screen.setParent(cabinet)
[1854]              // Position different parts of the arcade cabinet, etc
[1855]
[1856]              // Add the cabinet entity to the RealityView.
[1857]              content.add(cabinet)
[1858]          } update: { content in
[1859]              joystick.updateRotationWithDragLocation(dragLocation))
[1860]          }
[1861]          // Create a `DragGesture` that only triggers if the `joystick` entity
[1862]          // is targeted.
```

```

[1863]         .gesture(DragGesture(coordinateSpace: .global)
[1864]         .targetedToEntity(joystick)
[1865]         .updating($dragLocation) { value, state, _ in
[1866]             state = value.convert(value.location3D,
[1867]                                     from: .global, to: joystick.parent)
[1868]         }
[1869]     )
[1870]     .gesture(MagnifyGesture()
[1871]     .targetedToEntity()
[1872]     .updating($inProgressScaleFactor) { value, state, _ in
[1873]         state = value.magnification
[1874]     }.onEnded { value in
[1875]         cabinetScale *= inProgressScaleFactor
[1876]     }
[1877] )
[1878] }
[1879] }
[1880] ```

```

[1881]

[1882] Because we're using the 2D gesture system, the two `gesture` modifiers create exclusive gesture relationships by default. If the user drags on the joystick entity, the `DragGesture` will activate and cause `dragLocation` to update, in turn causing the `update` block

[1883] of the `View` to run. During the drag, the `MagnifyGesture` would not be allowed to begin. On the other hand if the user dragged on an entity other than the joystick, no gesture would trigger until a second hand starts dragging, at which point the `MagnifyGesture` would trigger, resulting in scaling the arcade cabinet entity.

[1884] In addition to targeting specific entities, a `targetedToEntity(where: QueryPredicate<Entity>)` which filters to only entities that are included in the results of the QueryPredicate passed in.

[1885] myView.gesture(RotateGesture3D().targetedToEntity(where:  
 .has(MyRotatableComponent.self))

[1886]

[1887] // A custom component to store some kind of rotation, ostensibly applied

[1888] // to the entity in a custom system.

[1889] struct MyRotatableComponent: Component {

[1890] var rotation: Rotation3D

[1891] var isEnabled: Bool

[1892] ...

[1893] }

[1894] ``

[1895]

[1896] In this example only entities with the `MyRotatableComponent`, which is a developer-defined

[1897] custom component, can be targeted by the `RotateGesture3D`.

[1898]

[1899] extension Gesture {

- [1900]       /// Require this gesture to target an entity.
- [1901]       ///
- [1902]       /// - Returns: A `RealityCoordinateSpaceConverting` value containing
- [1903]       /// the original gesture value along with the targeted entity.
- [1904]       public func targetedToAnyEntity() -> some  
Gesture<EntityTargetValue<Self.Value>>
- [1905]
- [1906]       /// Require this gesture to target `entity` or a descendant of `entity`.
- [1907]       ///
- [1908]       /// - Parameter entity: The entity the gesture should target.
- [1909]       ///
- [1910]       /// - Returns: A `RealityCoordinateSpaceConverting` value containing
- [1911]       /// the original gesture value along with the targeted entity.
- [1912]       public func targetedToEntity(\_ entity: Entity) -> some  
Gesture<EntityTargetValue<Self.Value>>
- [1913]
- [1914]       /// Require this gesture to target an entity that can be found in the results of the  
specified QueryPredicate
- [1915]       ///
- [1916]       /// - Parameter query: a `QueryPredicate<Entity>` to filter which entity the  
gesture should target
- [1917]       ///
- [1918]       /// - Returns: A `RealityCoordinateSpaceConverting` value containing

[1919]       /// the original gesture value along with the targeted entity.

[1920]       public func targetedToEntity(where query: QueryPredicate<Entity>) -> some  
Gesture<EntityTargetValue<Self.Value>>

[1921]       }

[1922]

[1923]       /// A value containing an original gesture value along with a targeted entity.

[1924]       /// Spatial data from `Value` can be converted to and from the entity

[1925]       /// using functions defined in `RealityCoordinateSpaceConverting`.

[1926]       ///

[1927]       /// For example, to convert a `DragGesture`'s `location3D` to `entity`'s

[1928]       /// parent:

[1929]       ///

[1930]       /// DragGesture(coordinateSpace: .global).targetedToEntity().updating(\$state) {  
state, value, \_ in

[1931]       ///     let location = value.convert(  
[1932]       ///         value.location3D, from: .global, to: value.entity.parent  
[1933]       ///     )  
[1934]       ///     ...  
[1935]       ///     }

[1936]       @dynamicMemberLookup

[1937]       public struct EntityTargetValue<Value>: RealityCoordinateSpaceConverting {

[1938]       /// The gesture value updated by the gesture.

[1939]       public var var gestureValue: Value

[1940]

[1941]       /// The targeted entity.

[1942]       public var entity: Entity

[1943]

[1944]       subscript<T>(dynamicMember keyPath: KeyPath<Value, T>) -> T

[1945]       }

[1946]

[1947]       extension EntityTargetValue: Equatable where Value: Equatable {

[1948]       public static func == (lhs: EntityTargetValue<Value>,

[1949]                               rhs: EntityTargetValue<Value>) -> Bool

[1950]       }

[1951]       ``

[1952]

[1953]

[1954]       For support on other platforms, the `targetedToEntity` modifier can provide similar functionality in `View` as it does on OS. However, because gestures on those platforms are fundamentally 2D, it's likely that we would want `EntityTargetValue` to assist with converting spatial data differently, for example by projecting 2D locations into the 3D scene.

[1955]       For this, we could have a `RealityCoordinateSpaceProjectable` protocol and have `EntityTargetValue` conform to that on 2D platforms.

[1956]       protocol RealityCoordinateSpaceProjectable {

```

[1957]     func project(
[1958]         _ point: CGPoint, from space: CoordinateSpace,
[1959]         to entity: Entity? = nil
[1960]     ) -> SIMD3<Float>
[1961]
[1962]     func unproject(
[1963]         _ point: SIMD3D<Float>, from entity: Entity? = nil,
[1964]         to space: CoordinateSpace
[1965]     ) -> CGPoint
[1966] }
[1967] ```

[1968] In the original version of this proposal, `targetedToEntity` had two variants, a
version that took an `Entity?`

[1969] and a version that took a particular `Component.Type` like so:

[1970]
[1971] extension Gesture {
[1972]     /// Require this gesture to target an entity, or a descendant of an entity,
[1973]     /// in order to activate.
[1974]     ///
[1975]     /// - Parameter entity: An entity that must be targeted by the gesture
[1976]     /// in order to activate. If `nil` is specified, then the gesture
[1977]     /// activates if it targets any `Entity`.

```

```

[1978]    ///

[1979]    /// - Returns: A `RealityCoordinateSpaceConverting` value containing

[1980]    /// the original gesture value along with the targeted entity.

[1981]    public func targetedToEntity(_ entity: Entity? = nil)

[1982]    -> some Gesture<EntityTargetValue<Self.Value>>

[1983]

[1984]    /// Require this gesture to target an entity that has the specified

[1985]    /// component type attached.

[1986]    ///

[1987]    /// - Parameter component: A component type used to filter the potential

[1988]    /// targets for this gesture.

[1989]    ///

[1990]    /// - Returns: A `RealityCoordinateSpaceConverting` value containing

[1991]    /// the original gesture value along with the targeted entity.

[1992]    public func targetedToEntity<RestrictedComponent: Component>(

[1993]        having component: RestrictedComponent.Type

[1994]    ) -> some Gesture<EntityTargetValue<Self.Value>>

[1995]    }

[1996]    ```

[1997]

[1998]    We determined it was confusing to have an optional `Entity` where `nil` means
    "any entity".

```

[1999] Rather than exposing a new variant of `targetedToEntity` like  
 `targetedToAnyEntity`, we think

[2000] the explicit modeling of the filter using the `QueryPredicate<Entity>` type makes  
 it easier to understand

[2001] the various possible filtering behaviors, especially for those already familiar with  
 3D framework.

[2002]

[2003] One alternative is to implement these behaviors using a new gesture called  
 `EntityGesture` whose value would be similar to `EntityTargetValue`. It was determined that  
 this version of the API is harder to read and easier to get wrong (it's rather easy to put the  
 `onChanged` on the `DragGesture` instead of the `EntityGesture<DragGesture>`.

[2004]

[2005] struct EntityGestureDemoView: View {

[2006] var body: some View {

[2007] View { content in

[2008] ...

[2009] } update: { \_ in }

[2010] .gesture(EntityGesture(DragGesture()).onChanged { event in

[2011] ...

[2012] })

[2013] }

[2014] }

[2015] ```

[2016] In some examples, since `QueryPredicate` already existed, it would be better for API continuity and cleanliness to use a system that was already in place, rather than building a new one. If we built out `GestureEntityTarget` it would have looked something like this:

[2017]

[2018] ```swift

[2019] /// A filter on a `Gesture`'s targeted `Entity`. If the filter

[2020] /// rejects a `Gesture`'s entity, then the `Gesture` does not activate and fails immediately.

[2021] public struct GestureEntityTarget {

[2022] /// A filter that requires a `Gesture` to target any `Entity`.

[2023] /// If the gesture does not target an `Entity`, the `Gesture` fails.

[2024] static var any: GestureEntityTarget { get }

[2025]

[2026] /// A filter that requires a `Gesture` to target `entity` or a descendant of `entity`.

[2027] static func descendingFrom(\_ entity: Entity) -> GestureEntityTarget

[2028]

[2029] /// A filter that requires a `Gesture` to target an `Entity` that has a component

[2030] /// of type `componentType`. If `includingAncestors` is true, then the `Gesture` may

[2031] /// activate if any ancestor `Entity` of the targeted `Entity` has a component of type `componentType`.

[2032] static func withComponent(\_ componentType: any Component.Type, includingAncestors: Bool = true) -> GestureEntityTarget

[2033]

```

[2034]    /// A filter that allows a `Gesture` to activate only if `filter` returns `true`
[2035]    /// given the `Entity` targeted by the `Gesture`.
[2036]    static func filter(_ filter: @escaping (Entity) -> Bool) -> GestureEntityTarget
[2037]    }
[2038]    ```
[2039]

```

[2040] There are a number of use cases for converting between 2D and 3D content in an app. For example, use of the new `GeometryReader3D` can inform a child `View` about the available space it has to configure its `Entity` hierarchy like so:

```

[2041]    GeometryReader3D { proxy3D in
[2042]        View { content in
[2043]            let frameInPoints = proxy3D.frame(in: .local)
[2044]            let sizedEntity = MyCoolModel(boundingBox: frameInPoints)
[2045]            content.add(sizedEntity)
[2046]        }
[2047]    }
[2048]    ```

```

[2049] In this example, `MyCoolModel` takes a size in order to figure out how large to make itself. However, the frame we get from `GeometryReader3D` is in points so that it can be useful in non-3D contexts as well. This leaves a gap in usefulness, as it's highly likely that `MyCoolModel` would prefer to deal with coordinates in the coordinate space of the `View`'s 3D content itself.

[2050] Additionally, 2D gestures on `View` content have similar needs. The new `targetedToEntity` Gesture modifier allows a 2D gesture to gate its activation on whether it

targets an entity. However, in order for the gesture to be useful in the context of `View`, the developer needs a way to convert spatial data into the coordinate space of an `Entity` inside the View.

[2051] This proposal aims to solve both these problems, among other similar ones related to conversions between 2D and 3D coordinate spaces.

[2052] In some examples, there are two new protocols, `RealityCoordinateSpace`, and `RealityCoordinateSpaceConverting`,

[2053] to make it easy to convert spatial data between 2D and 3D coordinate spaces. Conformance to `RealityCoordinateSpace` means that a `RealityCoordinateSpaceConverting` is able to convert properties to and from a `2D` coordinate space to the conformer. This protocol would have an underscored method so that we can add conformance to types internally and not allow conformance of arbitrary types.

[2054] `ViewInlineContent` would be an example of something that would conform so that the developer has a way to convert data to and from the root of their View.

[2055] Additionally, to handle immersive VR cases where the developer is dealing with content directly at the root of their volumetric window or `Stage`, a `RealityCoordinateSpace.scene` static property representing the `3D framework` `Scene`. This allows for easy conversions to and from a coordinate space aligned with ARKit's world anchor tracking and hand anchor APIs. Having `RealityCoordinateSpace` allows us to avoid ambiguity about what converting to and from `nil` means across different contexts.

[2056] In the proposal for "2D Gestures in View", using these new protocols with the new `targetedToEntity` Gesture modifier. For example, the developer could convert a 3D location from a `SpatialTapGesture` into a View's entity hierarchy like so:

[2057]

[2058] struct SpatialTapSphere: View {

[2059] var body: some View {

[2060] View { content in

```
[2061]         let sphere = ModelEntity(  
[2062]             mesh: .generateSphere(radius: 0.1),  
[2063]             materials: [SimpleMaterial(color: .blue, roughness: 0.25, isMetallic:  
false)],  
[2064]             collisionShape: .generateSphere(radius: 0.1),  
[2065]             mass: 1  
[2066]         )  
[2067]         sphere.physicsBody!.mode = .static  
[2068]         sphere.components.set(InputTargetComponent())  
[2069]         content.add(sphere)  
[2070]     }  
[2071]     .gesture(SpatialTapGesture().targetedToEntity().onEnded { value in  
[2072]  
[2073]         // Convert the drag location from 2D local space of the View  
[2074]         // to the targeted entity's local space.  
[2075]         let tapPoint = value.convert(  
[2076]             value.location3D,  
[2077]             from: .local,  
[2078]             to: value.entity  
[2079]         )  
[2080]  
[2081]         // Create a sphere at the tap location and add it as a child of the
```

```

[2082]         // targeted entity.

[2083]         let newEntity = ModelEntity(mesh: .generateSphere(radius: 0.01),

[2084]             materials: [SimpleMaterial(color: .red, roughness: 0.25, isMetallic:
false)])

[2085]         newEntity.position = tapPoint

[2086]         entity.addChild(newEntity)

[2087]     })

[2088] }

[2089] }

[2090] ```

[2091]

[2092] As mentioned in the Motivation section, this protocol would make
`GeometryReader3D` useful for determining how much space a `View` has to lay out its 3D
content:

[2093] ```swift

[2094] GeometryReader3D { proxy3D in

[2095]     RealityView { content in

[2096]         // Get the 3D frame size in points

[2097]         let frameInPoints = proxy3D.frame(in: .local)

[2098]         let boundingBoxInMeters = content.convert(sizeInPoints, from: .local, to:
content)

[2099]         let sizedEntity = MyCoolModel(boundingBox: boundingBoxInMeters)

[2100]         content.add(sizedEntity)

```

```
[2101]     }
[2102] }
[2103] ```
[2104] ```swift
[2105] /// A 3D coordinate space that exists within a hierarchy.
[2106] /// Any `RealityCoordinateSpaceConverting` can convert spatial data between
[2107] /// 2D `CoordinateSpace`s and `RealityCoordinateSpace`s.
[2108] public protocol RealityCoordinateSpace {
[2109]     // Internal implementation, making this protocol not publicly conformable.
[2110]     func _resolve(in context: _RealityCoordinateSpaceContext) ->
        _ResolvedRealityCoordinateSpace
[2111] }
[2112]
[2113] // Internal coordinate space implementation details.
[2114] public struct _ResolvedRealityCoordinateSpace {}
[2115]
[2116] // Internal coordinate space implementation details.
[2117] public struct _RealityCoordinateSpaceContext {}
[2118]
[2119] public extension Entity: RealityCoordinateSpace {}
[2120]
[2121] /// The root coordinate space of the 3D.
```

[2122]     public extension RealityViewInlineContent: RealityCoordinateSpace,  
RealityCoordinateSpaceConverting {}

[2123]

[2124]     /// Root coordinate space of the 3D `Scene`.

[2125]     extension RealityCoordinateSpace where Self == SceneCoordinateSpace {

[2126]         public static var scene: Self

[2127]     }

[2128]

[2129]     /// The `RealityCoordinateSpace` representing the root `Scene` coordinate space.

[2130]     public struct SceneCoordinateSpace: RealityCoordinateSpace {}

[2131]

[2132]     /// A value that can be converted between 2D `CoordinateSpace`s and

[2133]     public protocol RealityCoordinateSpaceConverting {

[2134]

[2135]         /// Convert a 3D point from a 2D `CoordinateSpace` to an entity.

[2136]         func convert(  
[2137]             \_point: Point3D,  
[2138]             from space: CoordinateSpace,  
[2139]             to realitySpace: some RealityCoordinateSpace  
[2140]         ) -> SIMD3<Float>

[2141]

[2142]         /// Convert a 3D point from an entity to a 2D `CoordinateSpace`.

```
[2143]     func convert(  
[2144]         _ point: SIMD3<Float>,  
[2145]         from realitySpace: some RealityCoordinateSpace,  
[2146]         to space: CoordinateSpace  
[2147]     ) -> Point3D  
[2148]  
[2149]     /// Convert a 3D rect from a 2D `CoordinateSpace` to an entity.  
[2150]     func convert(  
[2151]         _ rect: Rect3D,  
[2152]         from space: CoordinateSpace,  
[2153]         to realitySpace: some RealityCoordinateSpace  
[2154]     ) -> BoundingBox  
[2155]  
[2156]     /// Convert a `BoundingBox` from an entity to a 2D `CoordinateSpace`.  
[2157]     func convert(  
[2158]         _ boundingBox: BoundingBox,  
[2159]         from realitySpace: some RealityCoordinateSpace,  
[2160]         to space: CoordinateSpace  
[2161]     ) -> Rect3D  
[2162]  
[2163]     /// Convert a 3D transform from a 2D `CoordinateSpace` to an entity.
```

```

[2164]     func convert(
[2165]         _ transform: AffineTransform3D,
[2166]         from space: CoordinateSpace,
[2167]         to realitySpace: some RealityCoordinateSpace
[2168]     ) -> Transform
[2169]
[2170]     /// Convert a 3D transform from an entity to a 2D `CoordinateSpace`.
[2171]     func convert(
[2172]         _ transform: Transform,
[2173]         from realitySpace: some RealityCoordinateSpace,
[2174]         to space: CoordinateSpace
[2175]     ) -> AffineTransform3D
[2176] }
[2177] ```

```

[2178] On other platforms, spatial data is fundamentally 2D. For both the `targetedToEntity` and `View` use cases, we'd likely want a different way to convert spatial data to and from 3D framework. For example, one could expose a `RealityCoordinateSpaceProjectable` that could project and unproject 2D data

```

[2179] protocol RealityCoordinateSpaceProjectable {
[2180]     func project(
[2181]         _ point: CGPoint, from space: CoordinateSpace,
[2182]         to realityCoordinateSpace: some RealityCoordinateSpace
[2183]     ) -> SIMD3<Float>

```

[2184]

[2185]       func unproject(

[2186]            \_point: SIMD3D<Float>,

[2187]            from realityCoordinateSpace: some RealityCoordinateSpace,

[2188]            to space: CoordinateSpace

[2189]        ) -> CGPoint

[2190]        }

[2191]        '''

[2192]       To help with consistency, one could replace the existing conversion functions between `Entity` instances with functions defined by `RealityCoordinateSpaceConverting`. This might also resolve ambiguity about what `nil` means in those conversion functions.

[2193]       Without a static `.scene` property, it would likely still be possible to convert values to and from a "root" scene space, but it would be more difficult, requiring the developer to create an `Entity` that tracks the identity transform relative to its place in the hierarchy, which could be invalidated if the `View` containing it moves.

[2194]       Without a unified solution for 3D layout, and `targetedToEntity`, each of these solutions would likely need to expose their own APIs for similar kinds of conversions between 2D and 3D contexts. This will likely be overwhelming to developers.

[2195]       As described above, one aspect of the present technology is the gathering and use of data available from specific and legitimate sources for providing content. The present disclosure contemplates that in some instances, this gathered data may include personal information data that uniquely identifies or can be used to identify a specific person. Such personal information data can include audio data, voice data, demographic data, location-based data, online identifiers, telephone numbers, email addresses, home addresses, encryption information, data or records relating to a user's health or level of fitness (e.g., vital signs measurements, medication information, exercise information), date of birth, or any other personal information.

**[2196]** The present disclosure recognizes that the use of personal information data, in the present technology, can be used to the benefit of users. For example, the personal information data can be used for providing content.

**[2197]** The present disclosure contemplates that those 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 would be expected to implement and consistently apply privacy practices that are generally recognized as meeting or exceeding industry or governmental requirements for maintaining the privacy of users. Such information regarding the use of personal data should be prominently and easily accessible by users, and should be updated as the collection and/or use of data changes. Personal information from users should be collected for legitimate uses only. Further, such collection/sharing should occur only after receiving the consent of the users or other legitimate basis specified in applicable law. Additionally, such entities should consider taking 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. In addition, policies and practices should be adapted for the particular types of personal information data being collected and/or accessed and adapted to applicable laws and standards, including jurisdiction-specific considerations which may serve to impose a higher standard. For instance, in the US, collection of or access to certain health data may be governed by federal and/or state laws, such as the Health Insurance Portability and Accountability Act (HIPAA); whereas health data in other countries may be subject to other regulations and policies and should be handled accordingly.

**[2198]** Despite the foregoing, the present disclosure also contemplates embodiments 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 techniques for rendering content, the present technology can be configured to allow users to select to “opt in” or “opt out” of participation in the collection and/or sharing of personal information data during registration for services or anytime thereafter. In addition to providing “opt in” and “opt out” options, the present disclosure contemplates providing

notifications relating to the access or use of personal information. For instance, a user may be notified upon downloading an app that their personal information data will be accessed and then reminded again just before personal information data is accessed by the app.

**[2199]** Moreover, it is the intent of the present disclosure that personal information data should be managed and handled in a way to minimize risks of unintentional or unauthorized access or use. Risk can be minimized by limiting the collection of data and deleting data once it is no longer needed. In addition, and when applicable, including in certain health related applications, data de-identification can be used to protect a user's privacy. De-identification may be facilitated, when appropriate, by removing identifiers, controlling the amount or specificity of data stored (e.g., collecting location data at city level rather than at an address level or at a scale that is insufficient for facial recognition), controlling how data is stored (e.g., aggregating data across users), and/or other methods such as differential privacy.

**[2200]** Therefore, although the present disclosure broadly covers use of personal information data to implement one or more various disclosed embodiments, the present disclosure also contemplates that the various embodiments can also be implemented without the need for accessing such personal information data. That is, the various embodiments of the present technology are not rendered inoperable due to the lack of all or a portion of such personal information data.

**[2201]** FIG. 8 illustrates an electronic system 800 with which some examples of the subject technology may be implemented. The electronic system 800 can be, and/or can be 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.

**[2202]** The bus 808 collectively represents all system, peripheral, and chipset buses that communicatively connect the numerous internal devices of the electronic system 800. In some examples, 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 examples.

**[2203]** 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 some examples, 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.

**[2204]** In some examples, 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 some examples, 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 some examples.

**[2205]** 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 800. 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 800. 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. Some examples may include

devices that function as both input and output devices, such as a touchscreen. In these examples, 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.

**[2206]** 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.

**[2207]** 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.

**[2208]** Some examples 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.

**[2209]** While the above discussion primarily refers to microprocessor or multi-core processors that execute software, some examples are performed by one or more integrated circuits, such as application specific integrated circuits (ASICs) or field programmable gate arrays (FPGAs). In some examples, such integrated circuits execute instructions that are stored on the circuit itself.

**[2210]** 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.

**[2211]** To provide for interaction with a user, some examples 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.

**[2212]** 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 example 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).

**[2213]** 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 embodiments, 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.

**[2214]** Some examples 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.

**[2215]** 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.

**[2216]** 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 some examples, the tangible computer-readable storage medium can be directly coupled to a computing device, while in other examples, 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.

**[2217]** 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.

**[2218]** While the above discussion primarily refers to microprocessor or multi-core processors that execute software, some examples are performed by one or more integrated circuits, such as ASICs or FPGAs. In some examples, such integrated circuits execute instructions that are stored on the circuit itself.

**[2219]** 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.

**[2220]** 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 some examples, multitasking and parallel processing may be

advantageous. Moreover, the separation of various system components in the examples described above should not be understood as requiring such separation in all examples, 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.

**[2221]** 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.

**[2222]** 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.

**[2223]** 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 some examples, 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.

**[2224]** 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 examples, some examples, other examples, such examples, one example, for example, 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.

**[2225]** 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 examples. 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.

**[2226]** 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”.

**[2227]** 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.

## CLAIMS

What is claimed is:

1. A method, comprising:
  - receiving, from an application, a request to add a two-dimensional entity at a first location in a three-dimensional environment, wherein the 3D environment includes one or more 3D entities; and
  - in response to receiving the request to add the 2D entity at the first location in the 3D environment:
    - adding a first 3D entity to the first location in the 3D environment;
    - rendering, via a 2D framework, a representation of the 2D entity; and
    - rendering, via a 3D framework, a representation of a second 3D entity by:
      - performing one or more operations on the representation of the 2D entity; and
      - placing the representation of the second 3D entity at the first location.
2. The method of claim 1, wherein an observable representation of the first 3D entity is not present in the 3D environment.
3. The method of any one of claims 1-2, wherein performing the one or more operations include rotating the representation of the 2D entity, shrinking the representation of the 2D entity, enlarging the representation of the 2D entity, stretching the representation of the 2D entity, modifying a shape of the representation of the 2D entity, modifying a color characteristic of the representation of the 2D entity, or one or more combinations thereof.
4. The method of any one of claims 1-3, further comprising:
  - sending, to a display process, a request to display a representation of the one or more 3D entities.
5. The method of any one of claims 1-4, further comprising:
  - after placing the representation of the second 3D entity at the first location, detecting movement of a visible representation of a third 3D entity of the one or more 3D entities from

a second location to a third location, wherein the third location is different from the second location, and wherein the second location is a first distance from the first location;

in response to detecting the movement of the visible representation of the third 3D entity, placing the representation of the second 3D entity at a fourth location different from the first location, wherein the fourth location is the first distance from the third location.

6. The method of any one of claims 1-5, further comprising:

receiving a change to the 2D entity; and

in response to receiving the change to the 2D entity, rendering, via the 3D framework, a second representation of the 2D entity, wherein the second representation is different from the representation of the 2D entity.

7. The method of any one of claims 1-6, further comprising:

detecting, via one or more input devices, an input directed to the representation of the second 3D entity; and

in response to detecting the input directed to the representation of the second 3D entity and in accordance with a determination that a first set of one or more criteria is satisfied:

translating the input to a second input different from the input; and

sending, to the application, an indication of the second input.

8. The method of any one of claims 1-7, further comprising:

detecting an interaction between the representation of the second 3D entity and a representation of a fourth 3D entity of the one or more 3D entities; and

in response to detecting the interaction:

rendering, via the 2D framework, an updated representation of the 2D entity;

and

rendering, via the 3D framework, an updated representation of the second 3D entity.

9. The method of any one of claims 1-8, further comprising:

receiving, from the application, information corresponding to the 2D entity;

after receiving the request to add the 2D entity at the first location in the 3D environment, rendering the 3D environment, wherein the information affects rendering a

visual representation of a fifth 3D entity in the 3D environment, and wherein the fifth 3D entity does not correspond to the 2D entity.

10. The method of any one of claims 1-9, wherein the first 3D entity is added to a graph of the 3D environment.
11. The method of any one of claims 1-10, wherein the first 3D entity is included in a second 3D environment different from the 3D environment.
12. The method of any one of claims 1-11, wherein the first 3D entity includes a size, a location, an orientation, a pose, or any combination thereof.
13. The method of any one of claims 1-12, further comprising:
  - receiving, from a second application different from the application, a request to add a sixth 3D entity to the 3D environment; and
  - after receiving the request to add the sixth 3D entity to the 3D environment, rendering, via the 3D framework, a representation of the sixth 3D entity.
14. The method of any one of claims 1-12, wherein the application causes the one or more 3D entities to be rendered for the 3D environment.
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 an application, a request to add a two-dimensional entity at a first location in a three-dimensional environment, wherein the 3D environment includes one or more 3D entities; and

in response to receiving the request to add the 2D entity at the first location in the 3D environment:

adding a first 3D entity to the first location in the 3D environment;

rendering, via a 2D framework, a representation of the 2D entity; and

rendering, via a 3D framework, a representation of a second 3D entity by:

performing one or more operations on the representation of the 2D entity; and

placing the representation of the second 3D entity at the first location.

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 an application, a request to add a two-dimensional entity at a first location in a three-dimensional environment, wherein the 3D environment includes one or more 3D entities; and

in response to receiving the request to add the 2D entity at the first location in the 3D environment:

adding a first 3D entity to the first location in the 3D environment;

rendering, via a 2D framework, a representation of the 2D entity; and

rendering, via a 3D framework, a representation of a second 3D entity by:

performing one or more operations on the representation of the 2D entity; and

placing the representation of the second 3D entity at the first location.

21. A computer system, comprising:

means for, receiving, from an application, a request to add a two-dimensional entity at a first location in a three-dimensional environment, wherein the 3D environment includes one or more 3D entities; and

in response to receiving the request to add the 2D entity at the first location in the 3D environment:

means for adding a first 3D entity to the first location in the 3D environment;

means for, rendering, via a 2D framework, a representation of the 2D entity;

and

rendering, via a 3D framework, a representation of a second 3D entity by:

means for performing one or more operations on the representation of the 2D entity; and

means for placing the representation of the second 3D entity at the first location.

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 an application, a request to add a two-dimensional entity at a first location in a three-dimensional environment, wherein the 3D environment includes one or more 3D entities; and

in response to receiving the request to add the 2D entity at the first location in the 3D environment:

adding a first 3D entity to the first location in the 3D environment;

rendering, via a 2D framework, a representation of the 2D entity; and

rendering, via a 3D framework, a representation of a second 3D entity by:

performing one or more operations on the representation of the 2D entity; and

placing the representation of the second 3D entity at the first location.

23. A method, comprising:

at a computer system in communication with one or more input devices:

detecting, via the one or more input devices, a first input corresponding to a three-dimensional environment;

in response to detecting the first input corresponding to the 3D environment:

in accordance with a determination that a first set of one or more criteria is satisfied, wherein the first set of one or more criteria includes a criterion that is satisfied when the input is directed to a first entity of a first type in the 3D environment:

translating the first input to a second input different from the first input; and

sending, to a first application, an indication of the second input;

and

in accordance with a determination that a second set of one or more criteria is satisfied, wherein the second set of one or more criteria includes a criterion that is satisfied when the input is directed to a second entity of a second type in the 3D environment, sending, to a second application, an indication of the first input, wherein the second type of entity is different from the first type of entity, and wherein the second set of one or more criteria is different from the first set of one or more criteria.

24. The method of claim 23, wherein an observable representation of the first entity of the first type is not present in the 3D environment.

25. The method of any one of claims 23-24, wherein an observable representation of the first entity of the second type is present in the 3D environment.

26. The method of any one of claims 23-25, further comprising:

in response to detecting the first input corresponding to the 3D environment and in accordance with a determination that a third set of one or more criteria is satisfied, wherein the third set of one or more criteria includes a criterion that is satisfied when the input is a first type of input directed to the first entity of the first type in the 3D environment, sending, to the first application, the indication of the first input, wherein the third set of one or more criteria is different from the first set of one or more criteria and the second set of one or more criteria, and wherein the first set of one or more criteria includes a criterion that is satisfied when the input is a second type of input different from the first type of input.

27. The method of any one of claims 23-26, wherein the first input includes a 3D gesture.
28. The method of claim 27, wherein the first input includes an air gesture.
29. The method of any one of claims 23-28, wherein the first input includes a gaze of a user directed to a location in the 3D environment.
30. The method of any one of claims 23-29, further comprising:  
in response to detecting the first input corresponding to the 3D environment:  
in accordance with a determination that a fourth set of one or more criteria is satisfied, wherein the fourth set of one or more criteria includes a criterion that is satisfied when the input is directed to a third entity of a third type in the 3D environment, sending, to a fourth application, an indication of a location corresponding to a first type of coordinate space for the first input, an orientation corresponding to the first type of coordinate space for the first input, a pose corresponding to the first type of coordinate space for the first input, a magnitude corresponding to the first type of coordinate space for the first input, or any combination thereof, wherein the first type of coordinate space is defined by the fourth application; and  
in accordance with a determination that a fifth set of one or more criteria is satisfied, wherein the fifth set of one or more criteria includes a criterion that is satisfied when the input is directed to a fourth entity of a fourth type in the 3D environment, sending, to the fourth application, an indication of a location corresponding to a second type of coordinate space for the first input, an orientation corresponding to the second type of coordinate space for the first input, a pose corresponding to the second type of coordinate space for the first input, a magnitude corresponding to the second type of coordinate space for the first input, or any combination thereof, wherein the second type of coordinate space is defined by the fourth application, wherein the fifth set of one or more criteria is different from the fourth set of one or more criteria, wherein the fourth type of entity is different from the third type of entity, and wherein the second type of coordinate space is different from the first type of coordinate space.
31. The method of any one of claims 23-30, wherein the second application is the first application.

32. The method of any one of claims 23-31, wherein the second input includes a two-dimensional representation of the first input.
33. The method of claim 32, wherein translating the first input to the second input includes modifying a representation of a respective input from having six degrees of freedom to two degrees of freedom.
34. A non-transitory computer-readable medium storing one or more programs configured to be executed by one or more processors of a computer system in communication with one or more input devices, the one or more programs including instructions for performing the method of any one of claims 23-33.
35. A computer system in communication with one or more input devices, 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 23-33.
36. A computer system in communication with one or more input devices, comprising:  
means for performing the method of any one of claims 23-33.
37. A computer program product, comprising one or more programs configured to be executed by one or more processors of a computer system in communication with one or more input devices, the one or more programs including instructions for performing the method of any one of claims 23-33.
38. 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 in communication with one or more input devices, the one or more programs including instructions for:  
detecting, via the one or more input devices, a first input corresponding to a three-dimensional environment;  
in response to detecting the first input corresponding to the 3D environment:

in accordance with a determination that a first set of one or more criteria is satisfied, wherein the first set of one or more criteria includes a criterion that is satisfied when the input is directed to a first entity of a first type in the 3D environment:

translating the first input to a second input different from the first input; and

sending, to a first application, an indication of the second input; and

in accordance with a determination that a second set of one or more criteria is satisfied, wherein the second set of one or more criteria includes a criterion that is satisfied when the input is directed to a second entity of a second type in the 3D environment, sending, to a second application, an indication of the first input, wherein the second type of entity is different from the first type of entity, and wherein the second set of one or more criteria is different from the first set of one or more criteria.

39. A computer system in communication with one or more input devices, 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:

detecting, via the one or more input devices, a first input corresponding to a three-dimensional environment;

in response to detecting the first input corresponding to the 3D environment:

in accordance with a determination that a first set of one or more criteria is satisfied, wherein the first set of one or more criteria includes a criterion that is satisfied when the input is directed to a first entity of a first type in the 3D environment:

translating the first input to a second input different from the first input; and

sending, to a first application, an indication of the second input;  
and

in accordance with a determination that a second set of one or more criteria is satisfied, wherein the second set of one or more criteria includes a criterion that is satisfied when the input is directed to a second entity of a second type in the 3D environment, sending, to a second application, an indication of the first input, wherein the second type of entity is different from the first type of entity, and wherein the second set of one or more criteria is different from the first set of one or more criteria.

40. A computer system in communication with one or more input devices, comprising:  
means for, detecting, via the one or more input devices, a first input corresponding to a three-dimensional environment;  
in response to detecting the first input corresponding to the 3D environment:  
in accordance with a determination that a first set of one or more criteria is satisfied, wherein the first set of one or more criteria includes a criterion that is satisfied when the input is directed to a first entity of a first type in the 3D environment:  
means for translating the first input to a second input different from the first input; and  
means for, sending, to a first application, an indication of the second input; and  
means for, in accordance with a determination that a second set of one or more criteria is satisfied, wherein the second set of one or more criteria includes a criterion that is satisfied when the input is directed to a second entity of a second type in the 3D environment, sending, to a second application, an indication of the first input, wherein the second type of entity is different from the first type of entity, and wherein the second set of one or more criteria is different from the first set of one or more criteria.
41. A computer program product, comprising one or more programs configured to be executed by one or more processors of a computer system in communication with one or more input devices, the one or more programs including instructions for:  
detecting, via the one or more input devices, a first input corresponding to a three-dimensional environment;  
in response to detecting the first input corresponding to the 3D environment:  
in accordance with a determination that a first set of one or more criteria is satisfied, wherein the first set of one or more criteria includes a criterion that is satisfied when the input is directed to a first entity of a first type in the 3D environment:  
translating the first input to a second input different from the first input; and  
sending, to a first application, an indication of the second input; and  
in accordance with a determination that a second set of one or more criteria is satisfied, wherein the second set of one or more criteria includes a criterion that is satisfied when the input is directed to a second entity of a second type in the 3D environment, sending, to a second application, an indication of the first input, wherein the second type of entity is

different from the first type of entity, and wherein the second set of one or more criteria is different from the first set of one or more criteria.

1/8







FIG. 1

2/8

**FIG. 2**

**3/8****FIG. 3**



**FIG. 4**






FIG. 5

6/8

**FIG. 6**

7/8

**FIG. 7**




FIG. 8

# INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2024/030181

## A. CLASSIFICATION OF SUBJECT MATTER

INV. G06F3/01 G06T19/00 G06T19/20  
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06T A63F G06F

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

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages                                                                                                                                                                                            | Relevant to claim No. |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| X         | US 2019/247130 A1 (STATE ANDREI [US] ET AL) 15 August 2019 (2019-08-15)<br>paragraphs [0004] - [0012], [0033]<br>paragraphs [0075] - [0088]; figure 3A<br>paragraphs [0119] - [0139]; figure 4A<br>paragraphs [0149] - [0155]; figure 7<br>paragraphs [0193], [0194]<br>----- | 1-6,8-22              |
| A         | US 2023/165639 A1 (DULIN ISAAC [US] ET AL) 1 June 2023 (2023-06-01)<br>paragraphs [0001] - [0007]<br>paragraphs [0043] - [0078]<br>-----                                                                                                                                      | 1-6,8-22              |
| A         | US 2009/123046 A1 (MIELEKAMP PETER [NL] ET AL) 14 May 2009 (2009-05-14)<br>paragraphs [0001] - [0012]<br>paragraphs [0027] - [0031]; figure 3A<br>-----                                                                                                                       | 1-6,8-22              |

☐ 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

21 August 2024

Date of mailing of the international search report

31/10/2024

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

Günther, Steffen

## INTERNATIONAL SEARCH REPORT

International application No.  
PCT/US2024/030181

### 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 - 6 , 8 - 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: 1-6, 8-22

Enhancing an extended reality scene by a new digital twin

---

2. claims: 7, 23-41

Correlating different user gestures with different kinds of virtual solids

---

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No  
PCT/US2024/030181

| Patent document<br>cited in search report | Publication<br>date | Patent family<br>member(s) | Publication<br>date |
|-------------------------------------------|---------------------|----------------------------|---------------------|
| US 2019247130 A1                          | 15-08-2019          | US 2019247130 A1           | 15-08-2019          |
|                                           |                     | US 2023233264 A1           | 27-07-2023          |
| -----                                     |                     |                            |                     |
| US 2023165639 A1                          | 01-06-2023          | US 2023165639 A1           | 01-06-2023          |
|                                           |                     | US 2023165640 A1           | 01-06-2023          |
| -----                                     |                     |                            |                     |
| US 2009123046 A1                          | 14-05-2009          | CN 101442934 A             | 27-05-2009          |
|                                           |                     | EP 2018119 A2              | 28-01-2009          |
|                                           |                     | JP 2009536543 A            | 15-10-2009          |
|                                           |                     | RU 2008148820 A            | 20-06-2010          |
|                                           |                     | US 2009123046 A1           | 14-05-2009          |
|                                           |                     | WO 2007132381 A2           | 22-11-2007          |
| -----                                     |                     |                            |                     |