Warren Robinett
Department of Computer Science
University of North Carolina
Chapel Hill, North Carolina
27599-3175

Synthetic Experience:
A Proposed Taxonomy

Abstract

A taxonomy is proposed to classify all varieties of technologically mediated experience. This includes virtual reality and teleoperation, and also earlier devices such as the microscope and telephone. The model of mediated interaction assumes a sensor–display link from the world to the human, and an action–actuator link going back from the human to the world, with the mediating technology transforming the transmitted experience in some way. The taxonomy is used to classify a number of example systems.

Two taxonomies proposed earlier are compared with the ideas presented in this paper. Then the long-term prospects of this field are speculated on, ignoring constraints of cost, effort, or time to develop. Finally, the ultimate limits of synthetic experience are discussed, which derive from properties of the physical universe and the human neural apparatus.

1 Introduction

The head-mounted display (HMD) has been used in two distinctly different kinds of applications: teleoperation, in which a human operator's senses are projected into a remote robot body, and virtual environments, in which the human can move through and interact with a three-dimensional computer-generated virtual world. New uses for the HMD are currently being discovered, such as in flight simulation, night vision goggles, microteleoperation, and augmented reality.

This paper proposes a taxonomy for classifying systems that incorporate an HMD. Systems are classified according to nine independent dimensions, each of which can take on a number of discrete values. The domain of this classification method is broad enough to also include technological precursors to the HMD, such as the telescope, microscope, television, and telephone.

This taxonomy attempts to impose some sense onto a very broad and very new area that is pregnant with unexplored possibilities. The method of extracting order from the chaos is to cleave the set of possible systems into a small number of disjoint sets by imposing distinctions. Each of these distinctions corresponds to a dimension of the taxonomy. For example, the dimension called "causality" distinguishes between teleoperation, in which the operator's actions affect the real world, and virtual environments, in which the operator's actions affect only a simulated world. Other dimensions have to do with the sensory modalities used by the systems (vision, hearing, touch, and others), the nature of the representations (or models) of the environments surrounding the user,

Presence, *Volume 1, Number 2, Spring 1992*

and displacements or scaling in time or space between the user's true position and the environment the user interacts with.

The attempt to classify, or even talk about, devices that produce reproductions of sensory experience immediately brings up many difficult issues in philosophy, psychology, and other fields. What is experience? What is reality? What is a representation of the world, or of an object? Is perfect reproduction of human sensory experience possible? Many more questions and issues of this sort could be listed here. Rather than be scared off by these difficult and complex issues, I have tried to mention the issues that I think are relevant to the discussion, and leave it to others to correct and clarify errors and omissions, if they desire. I hope this taxonomy can serve as a point of departure for us collectively to understand and develop head-mounted displays into useful tools.

The discussion in this paper is somewhat biased by my own experience in designing computer-simulated virtual worlds, and my lack of detailed knowledge of the work done over the last several decades in teleoperation, and perhaps other related fields. All of us in this diverse field have our specialties and our blind spots. My experience tells me that the distinctions that I put forward in this paper are important ones, and rather than waiting until I achieve broad knowledge of all the fields touching on virtual worlds, I put the ideas forward now to serve as a starting point for discussion.

The common theme of all these devices and systems is ***technologically mediated experience***. The older systems use optics or analog electronics to mediate and transform the user's experience, whereas the more recent systems rely heavily on computers and digital electronics. In both cases, the general model of technologically mediated experience is the same, as shown in Figure 1.

The new devices incorporating the HMD did not come out of nowhere, but are extensions and refinements of earlier devices and media. Media began to evolve thousands of years ago when prehistoric man created visual representations of the world using paint. Painting was followed by the telescope, the microscope, photography, the phonograph, the telephone, film, television, and video games. Each of these devices derives its usefulness from being able to modify, record, or transmit some aspect of human sensory experience. For each of these devices, sensory experience is captured, processed, and then displayed to a human user.

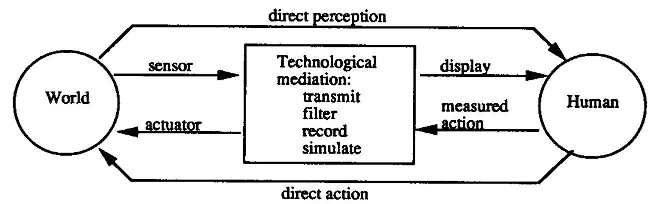

Figure 1. *Technologically mediated experience.*

The HMD is one further step along this evolutionary path. It improves on earlier visual media in being able to give the user a perception of a surrounding three-dimensional space, rather than just a look into a space from a fixed viewpoint. It is not simply a visual display technique, but rather a multisensory display technique (involving vision, the vestibular system, and the proprioceptive system) in which the visuals depicting the surrounding three-dimensional (3-D) virtual world are generated so as to match the user's voluntary head movements.

We offer some definitions:

- *natural experience:* directly perceiving the properties or behavior of something physically present before the perceiver.
- *synthetic experience:* perceiving a representation or simulacrum of something physically real rather than the thing itself.

The dictionary definitions (Webster's *Ninth New Collegiate Dictionary*) help to clarify the term "synthetic experience":

- **experience**
  1a: Direct observation of or participation in events as a basis of knowledge.
  3a: The conscious events that make up an individual's life.
  5: The act or process of directly perceiving events or reality.
- **synthetic**
  4b: Devised, arranged, or fabricated for special situations to imitate or replace usual realities.

5: Something resulting from synthesis rather than occurring naturally.

- **synthesis**
  1a: The composition or combination of parts or elements so as to form a whole.

The term *synthetic experience* encompasses virtual environments, teleoperation, other uses of the HMD, film, the telephone, video games, and most earlier media. It is meant to be synonymous with the term *technologically mediated experience,* used earlier. We limit the scope of synthetic experience to reproductions of sensory experience. We exclude verbal descriptions such as novels and oral story-telling.

We also exclude theatre from this classification system, though there is clearly a common thread running from story-telling to theatre to film. That theatre is in one sense a natural experience of watching human actors, and at the same time a recreation of a (hypothetical) earlier action, shows that it may be difficult to draw a distinct boundary between natural and synthetic experience. In the broadest sense, a student's mimicking of a tennis pro's serve is a reproduction of an earlier action, and role-playing in group therapy is a simulated experience. However, we limit the scope of synthetic experience to technologically mediated reproductions of sensory experience.

1.1 Examples of Synthetic Experience

Some of the most important current applications of the HMD are landmarks that help to map out the scope of synthetic experience. These examples are meant to illustrate the breadth of synthetic experience and are not meant to be definitions.

Virtual reality uses a stereoscopic, wide-angle HMD to create the illusion of a 3-D surrounding fantasy world, a 3-D video game that allows one or more players to get inside and interact with one another (Blanchard, Burgess, Harvill, Lanier, Lasko, Oberman, & Teilel, 1990).

Flight simulation also defines a simulated 3-D world in which actions have effects, but in this case the simulation is intended to accurately model the behavior of a real aircraft so as to give the pilot experience in dangerous situations without mistakes being fatal.

Teleoperation uses devices such as an HMD and force-feedback handgrip that are electronically linked to a distant robot body with a robot arm and a pair of video cameras on the robot head. The robot head turns to mimic the operator's head motions and the robot arm mimics hand motions, so that the operator's eyes and hands are effectively projected into the remote environment, and the operator can look around and do things through the robot body. The remote environment may be a dangerous one, such as the bottom of the ocean, inside a nuclear power plant, or in space.

Microteleoperation replaces the human-scale anthropomorphic robot of ordinary teleoperation with a microscope and micromanipulator, so as to give the operator the sense of presence and the ability to act in the microscopic environment. The scanning-tunneling microscope (STM) is well suited to microteleoperation since it uses a tiny probe scanned over the sample surface to capture a 3-D image of the surface (at atomic resolution), and the probe tip can also be used as a micromanipulator to interact with the sample material (Robinett, Taylor, Chi, Wright, Brooks, Williams, & Snyder, 1992). Images derived from this system are shown on the front and back cover of this issue.

Telecommunication is familiar to us through daily use of the telephone, and video teleconferencing extends this remote communication with other human beings to include the sense of sight, and to allow communication among groups of people rather than just two at a time. The operator of a tele-controlled robot is able to speak and listen to another human being in front of the robot.

Technological masquerade has been used to study intraspecies communication in animals by using recordings and sophisticated puppets to fool the animals into behaving as if they were interacting with another member of their species. This has been done extensively with recorded bird calls (Bright, 1984), and also with a computer-controlled robot bee, which was able to direct real bees to specific locations far from the hive by moving in the patterned "dance" that bees use to communicate and then dispensing a sweet liquid "sample" of the (pretended) distant pollen (Weiss, 1989). The connection with synthetic experience is that a human could potentially teleoperate a robot bee to attempt to communicate in real-time with real bees.

Augmented reality uses a see-through HMD, in which half-silvered mirrors allow the user to see through directly to the real world, and at the same time spatially superimpose the virtual world on top of the real world. The superimposed virtual world may be labels or diagrams located at specific points in the real world (Caudell & Mizell, 1992). It may also be information derived from sensors that is superimposed onto the user's direct view of the real world, as with, for example, helicopter pilots flying at night through canyons using forward-looking infrared (FLIR) sensors, who also have a direct view out into the darkness in case there is anything bright enough to see.

A *synthetic sense* is created when a sensor for a phenomenon that is imperceptible to human senses is linked to a display device. This gives the user the ability to perceive phenomena that are invisible, silent, and intangible without technological augmentation. Night vision goggles are an example of this. Another example, currently being prototyped, allows an obstetrician to use a see-through HMD to view data from a hand-held ultrasound scanner. The doctor can see and touch the abdomen of a pregnant woman, and sees the data from the ultrasound scanner superimposed at the location from which it came, giving the perception of seeing into the living tissue (Robinett, 1991a).

A *sensory prosthesis* corrects, amplifies, or otherwise improves the fidelity of an ordinary "built-in" human sense. Examples are corrective spectacles, sunglasses, and hearing aids. For people with defective or nonfunctional senses, *sensory substitution* can compensate for the disability. For example, the Opticon (Linvill, 1973) is an optical-to-tactile transducer array that allows blind people, after some training, to read from ordinary printed books by, in effect, running their finger over the printed text and feeling the black marks as raised bumps.

2 Dimensions of Synthetic Experience

The proposed taxonomy for classifying types of synthetic experience is shown in Table 1. The nine dimensions of the classification system are largely independent of one another, so the space of all possible types of synthetic experience should be conceived as a matrix (with nine dimensions) rather than a hierarchy. The first five dimensions describe the basic nature of the technological mediation in a synthetic experience device, whereas the last four dimensions have to do with which sensory channels and motor channels are employed.

2.1 Causality

The first dimension of the classification system, *causality,* makes the most fundamental distinctions among types of synthetic experience. The three possibilities are to *transmit, record,* or *simulate* experience. These three categories correspond to the way that we experience the world—not only do we experience the present, but we also remember the past and imagine the future. Replaying a recording has similarities with remembering: it is reexperiencing past events. Participating in an interactive simulation has similarities with imagining: it is trying out courses of action on an imaginary stage, perhaps to see what the consequences might be. Engaging in real-time transmitted experience through, for example, a teleoperator system, has similarities with normal active experience in the present: your actions affect the world (Robinett, 1991b).

The effect of voluntary actions is different in each of these cases. In a simulated virtual world (for example, in a flight simulator), actions have effects within that simulated world, but not in the real world. (There is no plane to crash; nobody will die.) In a virtual world which is a real-time reproduction of some part of the real world (for example, a pilot flying a remote-piloted aircraft), actions do affect the real world. (The plane can crash and burn.) In a recording of past events (for example, the "black box" recording of what happened in an airliner crash), what happened was recorded and actions by the user cannot change what happened. This dimension is called "causality" because, for the three cases of simulated, transmitted, and recorded experience, actions by the user can cause effects in the simulated world, effects in the real world, or no effects at all.

This dimension might possibly have been called "time," since recordings replay past events, transmitted experience takes place in real-time in the present, and

Table 1. *Classification System for Types of Synthetic Experience*

Dimension	Possibilities	Examples
Causality	Simulated	Flight simulator
	Recorded	Film
	Transmitted	Teleoperation
Model source	Scanned	Night vision goggles
	Constructed	Video game
	Computed	Computational fluid dynamics
	Edited	Film
Time	1-to-1	Film
	Accelerated (or retarded)	Time-lapse photography
	Frozen	Photograph
	Distorted	Edited video recording of event
Space	Registered	Night vision goggles
	Remote	Teleoperation
	Miniaturized (or enlarged)	Microteleoperation (STM)
	Distorted	STM with heights exaggerated
Superposition	Merged	Augmented reality
	Isolated	Virtual reality
Display type	HMD	Virtual reality
	Screen	Video game
	Speaker	Recorded music
	(Many more—see Table 3)	
Sensor type	Photomultiplier	Night vision goggles
	STM	Microteleoperation
	Ultrasound scanner	Medical "X-ray vision"
	(Many more—see Table 4)	
Action measurement type	Tracker and glove	Virtual reality
	Joystick	Video game
	Force feedback arm	Teleoperation
	(Many more—see Table 5)	
Actuator type	Robot arm	Teleoperation
	STM tip	Microteleoperation
	Aircraft flaps	Remote piloted aircraft
	(Many more—see Table 6)	

simulations sometimes are used to predict future events. However, simulations are not necessarily of the future (for example, a simulation of continental drift), so it is best to name this dimension "causality" to capture the real differences between transmitted, recorded, and simulated experience.

Figure 2 shows diagrams of the primary data flow for transmitted, recorded, and simulated experience, with

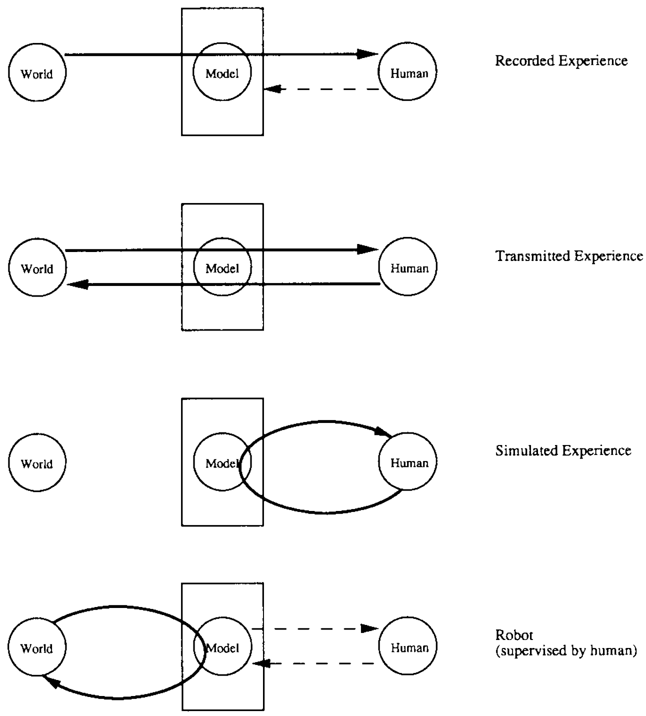

Figure 2. *Data flow for types of mediated experience.*

each shown as a special case of the diagram for technologically mediated experience in Figure 1. In the case of transmitted experience, the diagram of Figure 1 is a good model for the data flow—the user observes the world through the sensor–display data path and performs actions that affect the world through the action–actuator data path.

In the case of recorded experience, the sensor data are stored in some kind of memory device (such as magnetic tape), and at a later time these data are replayed through the display to the user. An actuator is not needed in this activity, and user actions are needed only to control the replay process itself.

In the case of simulated experience, the primary data path is from the measured actions of the user, through the simulation, and back through the display to the user. Again, the actuator is not needed, and the sensor channel is needed only if the simulated virtual world is based at least partly on scanned-in data from the real world.

The fourth diagram shown in Figure 2 is a variety of transmitted experience, with a data path introduced to allow autonomous actions by a telerobot, under supervision of a human operator. An operator could alternate between passive real-time observation of the telerobot's actions, and taking direct control of the telerobot's actions as in normal transmitted experience.

In a system in which all of these data paths are present, all of these modes of operation are possible. In the UNC Nanomanipulator Project (Robinett et al., 1992), in which a HMD and force-feedback arm control an STM, transmitted experience, recorded experience, simulated experience, and supervisory control are all possible. The user may directly control the STM tip through the force-feedback arm and modify the sample surface (transmitted experience). The user may record a snapshot or sequence of images of the surface and view them through the HMD at a later time (recorded experience). The user may manipulate simulated molecules through the force-feedback arm and HMD with no connection to the microscope (simulated experience). We plan later to allow the user to initiate algorithmically controlled modifications of the sample surface, with the possibility of intervening (transmitted experience with supervisory control).

2.2 Model Source

In a synthetic experience, the human user perceives a virtual world that is defined by a (possibly changing) database called a *model.* This model is stored, at least transitorily, in some kind of memory device. The model defines what the virtual world looks like, sounds like, and feels like, according to which display devices are available.

There are three main sources for this model data. A sensor can *scan* the real world to produce a model for later display to the user, a human artist or craftsman can laboriously *construct* a model, piece by piece, or a dynamic model can be *computed* on the fly by a computational model. Examples of these three cases are: live television, with the world scanned by the video camera; Disney-style animated cartoons, with each animation frame drawn by an artist; and computational fluid dynamics, where the simulation code generates new model data as needed. However, these cases are not exclusive, and a scanned-in model can be chopped up and *edited* to

Table 2. *Relative Scale and Displacement in Time and Space for Sensor and Display*

	Time	Space
Aligned	Transmitted in real-time, 1-to-1 time scale Live television	Registered, 1-to-1 scale Night vision goggles
Displaced	Recorded earlier, 1-to-1 time scale TV rebroadcast of live event	Remote, 1-to-1 scale Teleoperation
Scaled	Recorded earlier, accelerated (or retarded) time Slow motion instant replay on TV	Remote, expanded (or miniaturized) Microteleoperation
Distorted	Recorded earlier, distorted time TV event with dull parts edited out	Remote, distorted space Microteleoperation, exaggerated height

construct a model that is partly based on the real world, but is different. A good example of this is film, in which raw footage from the initial shooting is heavily edited, and some animated special effects are thrown in, to produce the final movie.

2.3 Time and Space

For data scanned in from the real world, in some cases (such as night vision goggles) the data will be displayed in exactly the location from which it was derived, whereas in other cases (such as teleoperation) the scan space is displaced from the display space. The display space may also differ in scale from the scan space, as in microteleoperation. The mapping from scan space to display space may include a spatial deformation.

Furthermore, the scan and display may be aligned or displaced in time (transmitted experience versus recorded experience). Scan time and display time may also differ in time scale, as with time-lapse photography. Display time could be related to scan time by a nonlinear distortion mapping, for example, in the replay of an explosion where initial events occur more rapidly than later events. Distorted time modeling has been used by researchers in telerobotics.

These possibilities may be summed up by saying that, for both time and space, the scan and display may be either aligned, displaced, differ in scale, or be related by a distortion mapping, as shown in Table 2. The relationship of Table 2 to the overall taxonomy of Table 1 is that Table 2 emphasizes the similarity of the values that can be assigned to the two dimensions "time" and "space" of the taxonomy.

Since this is a comparison of the time and space coordinates of the sensor and display, this comparison makes sense only when there *is* a sensor involved. A constructed model comes out of nothingness and therefore has no real world coordinates with which the display might align. Likewise, an edited model may have pieces that come from specific locations in the real world, but there is no way to match the whole of the model to the real world.

If we imagine two clocks displaying Greenwich Mean Time, one being scanned by the sensor and the other with the user beside the display, we may ask if the two clocks are displaying the same time and if they are running at the same rate. For transmitted experience, the two clocks must match from moment to moment, so 1-to-1 time scale is required. In replaying a recording, the clocks might run at the same rate but display different times. However, as with a VCR, the recording might also be played back in slow motion, faster than normal, in reverse, or paused with the action frozen. All these varieties of time progression are possible for any recording technique.

Since the distinction between transmitted and recorded experience is already covered by the causality dimension of the classification system, the time dimension of the classification system focuses on time scale, with the possibilities being *1-to-1* time-scale, *accelerated (or retarded)* time, *frozen* time, and *distorted* time.

In the same way that we used two clocks to judge the time-offset and time-scale differences between scan time and display time, we may also use two spatial markers to judge the offset and scale difference between scan space and display space. For this we use a pair of three-dimensional coordinate axes, one being scanned by the sensor and the other measuring the space the user occupies. For concreteness, let us imagine that the user wears a see-through HMD, which uses half-silvered mirrors to spatially superimpose the real and virtual worlds. In this case, the coordinate axes that are actually present in front of the user can be seen by the user with image of the scan-space coordinate axes optically superimposed. Both coordinate axes are ruled in centimeters. We can now ask whether the two axes are aligned or displaced from one another, and whether they appear to be the same or different sizes.

The main possibilities are that, relative to scan space, display space is *registered, displaced,* or *expanded* (*or miniaturized*). It is also possible to introduce various *distortion* mappings between scan space and display space. An example is in microteleoperation using the STM, where we wish to exaggerate the height variations of the sample being scanned so as to make very slight height steps more obvious.

2.4 Superposition

A virtual world may be *merged,* perhaps using half-silvered mirrors, with the surrounding real world. It may be convenient to combine the real and virtual worlds by using a video camera to capture an image of the real world and then doing a more sophisticated merge than is possible with optics. This example of using cameras to capture and merge the real world with a virtual one shows that the surrounding real world itself may be thought of as model, on an equal footing with the virtual world model, and the two models may be edited together, if desired. A powerful technique is to spatially superimpose two models of the same region of space, creating a sort of three-dimensional Rosetta Stone through the spatial correspondence of pairs of matching points in the two models.

Data from multiple sensors may be fused into a single virtual world model. An example is a conference telephone call. This may be thought of as automatic real-time editing, in which the data from three or more sensors (microphones) are integrated and then displayed through the speaker in each user's handset.

On the other hand, a display, particularly an HMD, may block out the real world and *isolate* the user within the virtual world.

2.5 Senses and Sensors

Display type and sensor type are two of the dimensions of the classification system, and a few examples are given in Table 1 where the classification system is defined. However, the complete range of displays can in principle cover every phenomenon that human beings possess sensory organs to detect. Likewise, the complete range of sensors encompasses all measurable or detectable phenomena. Table 3 lists human sensory channels and corresponding display devices, and Table 4 lists some sensors and the phenomena to which they are sensitive. Neither the list of display devices nor the list of sensors is exhaustive.

Many of the sensors listed detect phenomena that are imperceptible to human senses, and by linking such sensors to display devices, these imperceptible phenomena can be rendered visible, audible, touchable, or otherwise perceptible to a human being. A sensor-display linkage of this sort creates a synthetic sense, an apparatus that extends human perception and awareness (Robinett, 1991a).

Using a HMD as a display device offers the possibility of mapping sensed phenomena to specific locations relative to the body of the user. This is what vision and hearing do, and for those sensors that are able to establish the location or direction of the phenomena they sense, this positional information can be interpreted through the visual (and auditory) channels of the HMD to depict the sensor data as emanating from specific locations in space.

There are a considerable number of imperceptible

Table 3. *Human Sensory Channels and Display Devices*

Sensed phenomena	Sensory system	Display device
Visible light (400–700 nm)	Vision	Display screen (CRT, LCD, or other) Head-mounted display (HMD) Individual lights Dials and gauges
Vibrations in the air (20 Hz–20kHz)	Hearing	Speaker Headphones Headphones with spatialized sound
Force Vibration Surface texture Temperature	"Sense of touch"	Force-feedback device Buzzer Tactor array, air bladders Heater, cooler, fan
Chemical composition of air	Smell	Sensorama smell display
Chemical composition of food	Taste	
Acceleration of body	Vestibular system	Motion platform
Limb and body position	Proprioception	Exoskeleton with forced movements
Internal state of body (hunger, thirst, fatigue, etc.)	Interoceptors	Intravenous medical device to monitor and control contents of bloodstream
Damage to body	Pain	

phenomena. Every one of them can be given a visible form, or sound, or tactile representation. Every detectable phenomenon can be given a perceptible representation, regardless of its remoteness in space, time, scale, or time scale, and regardless of what form of energy or matter is being detected. By linking sensors and displays to create synthetic senses, every phenomenon that exists can be rendered directly perceptible to the human senses.

What do these imperceptible things look like? Since they are imperceptible, they do not look like anything. A representation must be invented, and choices present themselves. In general, many representations are possible for a given phenomenon, and different representations may be useful at different times. For example, a number of visual representations of molecules are used in different situations: touching sphere model, ball and stick model, solvent-accessible surface model, ribbon following the backbone of a protein. What these invisible things should look like is a graphic design problem that, in time, can be expected to settle out on the basis of informativeness, aesthetics, convention, and accidents of history.

2.6 Actions and Actuators

In the same way that sensors and human senses can be linked to cover all detectable phenomena, a linkage from manual and other input devices to actuators should be able to control any device or system designed to be controlled. This is a relatively unexplored area, with the main work so far having been done in teleoperation and remote piloted vehicles. Most other human tools, vehicles, environments, and instruments each have their own idiosyncratic locally operated control panels.

In a few years visual telepresence may be widely available, so that a person can move by virtual travel instantly

Table 4. *Sensors and What They Sense*

Sensed phenomena	Sensor
Visible light	Still camera
Visible light	Video camera
Sound	Microphone
Position of moving objects	Radar
Distance to object	Range-finder
Position and orientation of moving object	Tracker (6 DOF)
Inside of human body	Ultrasound scanner
	Computer-aided tomography (CAT) scan
	Nuclear Magnetic Resonance (NMR) imaging
Infrared light	Night vision goggles
Ultraviolet light	UV detector
X-rays	Fluoroscope
Magnetism	Electronic compass
Radiation	Geiger counter
3-D surface shape	3-D laser scanner
3-D topography of the earth	Aerial photography and photogrammetry
3-D surface of microscopic sample	Scanning-tunneling microscope
Image of distant object	Telescope
Chemical composition	Gas chromatograph
	Spectrograph
Movement and vibration	Accelerometer
	Gyroscope
Gravitation field variations	Mass detector

to distant locations, just as is now possible with the telephone for hearing only. If, at that time, most controllable devices are linked to the communications network, then it will be possible for a person to project by virtual travel to a distant location and initiate actions there through the actuators available at that site. For safety and security reasons, remote access will probably not be allowed for some types of devices, but for many devices

Table 5. *Human Motor Channels and Measurement Devices*

Motor channel	Behavior measurement device
Hands	Hand tracker (6 DOF)
	Hand-held pushbuttons
	Instrumented glove
	Keyboard
	Mouse
	Joystick
Feet	Foot pedal
Eyes (gaze-direction, blinking)	Gaze tracker
Head position	Head tracker (6 DOF)
Body posture	Instrumented body suit
Voice	Speech recognition
Breath	Breath controller
Heartbeat	EKG machine

Table 6. *Actuators Used in Teleoperated Systems*

Robot arm
STM tip
Remote piloted vehicle

it may make sense. Another issue is who has permission to control which devices. In spite of these probable limitations, we can still imagine a future world in which an enormous traffic of ghostly presences leap about the planet, manipulating distant parts of the world through briefly occupied robot bodies.

Table 5 lists human motor channels and some devices available to measure human actions. Table 6 lists some actuators that have so far been used in remote presence systems, and, of course, there are many more devices and systems that could potentially be controlled over the communication network.

2.7 Classification of Some Specific Systems

Table 7 shows the synthetic experience types of a number of specific systems and devices.

Table 7. *Examples of How Specific Systems Are Classified*

	Causality	Model source	Time	Space	Super-position	Display	Sensor	Action measure	Actuator
Teleoperation	Transmit	Scan	1-to-1	Remote	Isolated	HMD, force feedback arm	Camera on robot head	Force feedback arm	Robot arm
Microteleoperation	Transmit	Scan	1-to-1	Expanded	Isolated	HMD, force feedback arm	STM	Force feedback arm	STM tip
Remote piloted aircraft	Transmit	Scan	1-to-1	Remote	Isolated	Screen	Video camera	Joystick	Flap actuators in aircraft
Flight simulation	Simulate	Computed	1-to-1	Remote	Isolated	HMD, motion base	Satellite photography	Cockpit controls	—
Virtual reality game	Simulate	Construct	—	—	Isolated	HMD	—	Tracker, glove	—
Video game	Simulate	Construct	—	—	Isolated	Screen, speaker	—	Joystick	—
Augmented reality—helicopter	Transmit	Scan	1-to-1	Registered	Merge	HMD	FLIR	—	—
Night vision goggles	Transmit	Scan	1-to-1	Registered	Merge	HMD	Photomultiplier	—	—
Medical "X-ray vision"	Transmit	Scan	1-to-1	Registered	Merge	HMD	Ultrasound scanner	Hand tracker	—
Telephone	Transmit	Scan	1-to-1	Remote	Merge	Speaker	Microphone	Keypad	—
Live television	Transmit	Scan	1-to-1	Remote	—	Screen, speaker	Camera, microphone	—	—
Film	Record	Edit	1-to-1	Remote	—	Screen, speaker	Camera, microphone	—	—
Videocassette recorder	Record	Edit	1-to-1, fast, slow, frozen	Remote	—	Screen, speaker	Camera, microphone	Keypad	—
Time-lapse photography	Record	Scan	Accelerated	Remote	—	Screen	Camera	—	—
Photography	Record	Scan	Frozen	Remote	—	Print	Still camera	—	—
Painting	Record	Construct	Frozen	—	—	Canvas	—	—	—

Comparing different lines in the table suggests variations and extensions for some of the systems. For example, a hybrid of film and virtual reality would give us 3-D recording of earlier actions that the user could fly through to observe from any viewpoint.

A hybrid of microteleoperation with the STM and the video cassette recorder would allow rapid events occurring at the microscopic scale to be rapidly scanned as they occur, and then played back at a slower speed later, pausing and backing up to observe interesting events.

The dimensions of the classification system are largely independent of one another, so it is possible to go into the line for a given system and ask what kind of system would result by changing it along one dimension. The main dependencies between the dimensions are that transmitted experience requires 1-to-1 time scale, recorded experience needs no actuator or action measurement (except to control the replay), and simulated experience needs no actuator. Also, the dimensions of sensor-to-display relative time and space apply only to models scanned in by sensors from the real world.

3 Comparison with Earlier Taxonomies

This section discusses two broad taxonomies proposed earlier, which overlap with the taxonomy proposed in this paper. Then the taxonomies are compared with one another.

3.1 Realspace Imaging

In his paper "Elements of Realspace Imaging," Naimark proposed a taxonomy of six successively more complete methods for recording and reproducing experience (Naimark, 1991, 1992). These are monoscopic imaging, stereoscopic imaging, multiscopic imaging, panoramics, surrogate travel, and "real-time" imaging.

The first five categories cover techniques that allow successively greater freedom in achieving, during replay, arbitrary viewpoints in the 3-D scene that was recorded. He defines the last category, real-time imaging, as the "process of recording and displaying temporal sensory information indistinguishable from unmediated reality," in other words, high-fidelity, multisensory cinema. He gives a long list of techniques and problems in recording and replaying experience. Some of the issues and problems are orthoscopy, resolution, color, dynamic range, brightness, spatial consistency, accommodation, image stabilization, match-cuts, and storage capacity. Display techniques mentioned are mirrors, relief projection, holography, viewpoint-dependent imaging, and projection onto surrounding screens. Naimark primarily surveyed visual recording techniques in this paper, but he did also mention auditory, force, and vestibular (motion platform) displays.

3.2 AIP Cube

In his paper "Autonomy, Interaction, and Presence," Zeltzer proposed "a taxonomy of graphical simulation systems" (Zeltzer, 1992). The taxonomy consisted of three independent scalar dimensions that defined a space of possibilities, the "AIP cube." The dimension *autonomy* described the sophistication and dynamics of the model defining the virtual world. *Interaction* measured the degree to which user actions could affect what happened in the virtual world. *Presence* measured the sensory fidelity and breadth (the number of senses to which displays were aimed). The three dimensions were presented as rough lumped measures of the sophistication of the three key components of a simulated virtual world: the model defining the virtual world (autonomy), the input devices that let the user affect what happens in the virtual world (interaction), and the displays that let the user perceive the virtual world (presence).

Each dimension was measured by a scalar running from zero to one. Zeltzer commented that it was "not clear how to rigorously quantify" these dimensions, so assigning values for a given system would seem to be more a matter of judgment than measurement. He gave examples for each dimension. Autonomy ranged from a static model (autonomy = 0) to a fully autonomous agent (autonomy = 1). Interaction ranged from batch (interaction = 0) to real-time access to all model parameters (interaction = 1). Presence ranged from static graphics (presence = 0) to sensory stimulation indistinguishable from the real world (presence = 1).

Zeltzer gave examples of where specific systems fell within the AIP cube. He mapped the point (autonomy = 0, interaction = 0, presence = 0) to typical computing in the early 1960s: batch processing of simple graphic models with output on a plotter. The point (autonomy = 1, interaction = 1, presence = 1) mapped to "fully autonomous agents and objects which act and react according to the state of the simulation, and which are equally responsive to the actions of the human participant(s). In addition, sensory stimulation provided to the participant(s) in the virtual environment is indistinguishable from what would be expected in a physical setting." He stated that the (1,1,1) point was probably not achievable without direct neural connection. He described existing graphics systems that lacked one thing or another and therefore mapped to other corners of the cube. An interesting (hypothetical) system was "digital Shakespeare" at the point (autonomy = 1, interaction = 0, presence = 1). The user could view the action of the play from any viewpoint, and could rewind or fastforward, but would be unable to affect what happened in the play.

Zeltzer observed that "it is not possible to simulate the world in all its detail and complexity, so for a given task we need to carefully identify the sensory cues that must be provided in order for a human to accomplish the task, and match as closely as possible the human perceptual and motor performance required for the task." This technique is called *selective fidelity*.

3.3 Comparison with the Synthetic Experience Taxonomy

The domain of the realspace imaging taxonomy is limited to recordings of the real world, and thus it covers the same conceptual territory as recorded experience in the synthetic experience taxonomy. The distinctions made by the two taxonomies are independent of one another, and they neither contradict nor confirm one another.

Naimark's categories are monoscopic imaging, stereoscopic imaging, multiscopic imaging, panoramics, surrogate travel, and real-time imaging. These categories are arranged in a sequence that is pretty much the historical order of development, and they are also in order of increasing difficulty of implementation. In the synthetic experience taxonomy, the availability of stereoscopic HMDs with the ability to look around (panoramics) and to fly through the virtual world (surrogate travel) is assumed, and the distinctions made have to do with the displacements and distortions in time and space, the sensory modalities of the display and recording devices (sensors), and superposition of the real and virtual worlds.

One overall point of agreement between the two taxonomies is that an important area for research is this: three-dimensional scenes and actions could be recorded and then later replayed, perhaps through a head-mounted display, with the user allowed to move to any viewpoint within the 3-D scene. In this scenario, the user is not restricted to the viewpoint of the recording device that captured the scene or action, and the reproduction apparatus must synthesize the appropriate view from the user's position and the 3-D data originally recorded. Kanade's work is notable in this area (Gruss, Tada, & Kanade, 1992).

Zeltzer describes the AIP cube taxonomy as classifying graphical simulation systems, thus covering the same territory as the simulated experience subspace of the synthetic experience taxonomy. The "digital Shakespeare" example, in which the user can move to any viewpoint but cannot interact with the characters, comes rather close to the recorded experience scenario discussed above; but it is described as a simulation (with no interaction) rather than a recording of human actors.

There is an extremely loose correspondence between the three dimensions of the AIP cube and some dimensions of the synthetic experience taxonomy. The presence dimension of the AIP cube has to do with the degree of sensory coverage, and so maps to the display type dimension of synthetic experience. The interaction dimension of the AIP cube is somewhat related to which motor channels are available to control things in the simulation, as covered by the action measurement type dimension of synthetic experience. The autonomy dimension of the AIP cube has to with the nature of the model of the virtual world, as does the model source dimension of synthetic experience, but these two dimensions are focused on quite different aspects of the model: in the one case the model's responsiveness (autonomy), and in the other case the model's origin.

The AIP cube is more useful conceptually for suggesting untried possibilities than for actually classifying systems. The problem is that the possibilities lumped into each dimension do not map in any obvious fashion to a single linear scale. It is not clear how the autonomy, interaction, or presence of a system would be measured. In contrast, when using the synthetic experience taxonomy to classify a particular simulation system, it can be determined whether the model of the virtual world was scanned in, constructed, or computed; whether the virtual world is a miniaturized, distorted, displaced, accelerated, or frozen model of some part of the real world; whether the virtual world is visually merged with the real world; and which sensory and motor channels are employed.

The synthetic experience taxonomy attempts to cover a very broad intellectual territory. Its top-level distinction yields three categories: recorded experience, simulated experience, and transmitted experience. Realspace imaging and recorded experience cover the same ground, but lay out different boundaries within the territory. The AIP cube and simulated experience likewise cover the same ground, but with somewhat different internal boundaries. Transmitted experience covers the same ground as teleoperation.

4 Prospects and Limits of Synthetic Experience

The developments now underway in the various subclasses of synthetic experience are far from mature, and it seems clear that further exploitation of their inherent possibilities offers to humanity great increases in awareness, power, and the ability to effectively use our vast corpus of information and knowledge. With this power comes danger, and also the likelihood that these tools will become so pervasive and symbiotic with us that they change the very nature of being human. This is not necessarily a bad thing. Agriculture has changed humanity, and so have writing, mathematics, and science. But for an enterprise that portends great changes for humanity, it is sensible to ask: How far will it go? How far *can* it go? What is most likely to happen? What are the ultimate limits?

In this section of the paper, we will discuss some likely paths of development for various strands of synthetic experience, as best we can anticipate from our present knowledge of the physical universe and the human neural apparatus. We will try to extrapolate as far as we can, without regard to cost or effort or time to develop, but limited by what ultimately seems possible. Our model in this constrained speculation is Arthur C. Clarke's *Profiles of the Future* (Clarke, 1962).

We first imagine synthetic experience developed as far as we can foresee, and then consider the physical and psychological limits that form the ultimate boundary of the possible.

4.1 A Vision of the Potential Long-Term Development of Synthetic Experience

4.1.1 Perfect Fidelity of Synthetic Experience. Reproducing various aspects of experience has been the goal of most media from painting onward. The ultimate development of this aspect of synthetic experience is coverage of all human senses and perfect fidelity for each sensory channel, so that the human is unable to tell the difference between synthetic and natural experience. This criterion applies mainly to transmitted and simulated experience—you feel like you are really there. For recorded experience, since you have no ability to act, and must merely observe the action, you cannot be fooled into believing that you are having a natural experience.

4.1.2 Synthetic Senses Spanning All Detectable Phenomena. For every known detectable phenomenon involving energy or matter, and every kind of sensor that exists, a mapping may be made from the phenomenon to the built-in human senses. This implies that all these imperceptible phenomena have been given visual representations, or representations matched to other human senses.

4.1.3 Instant Travel at the Speed of Light. Transmitted experience permits experience at a distance, and with multiple teleoperated robots the human operator could switch his or her presence from one site to another as easily as people today call around the world to various sites on the telephone. With many sensors scanning the world in real-time, an integrated 3-D global database could be maintained in real time, so that the virtual travel from one site to another could be continuous motion with a changing viewpoint, rather than teleporting from site to site, as with the telephone. The speed of travel can be as fast as the user desires.

4.1.4 Apparent Magic. With all vehicles, tools, factories, libraries, and other controllable systems connected to the worldwide communications grid, a virtually present operator will be able to control devices or systems at any location. The input devices and control gestures will be arbitrary and independent of the devices being controlled. Arbitrary gestures activating, moving, creating, and destroying objects of the physical world will be similar in appearance and capabilities to the mythical idea of magic, for example, like the wizard Dr. Strange of the comic books.

4.1.5 Adventures in Microworlds. Microteleoperation will permit human operators to perceive and manipulate things in microscopic worlds ranging in scale from the merely tiny (where a bee is your size), down to the microscopic (where a bacterium is your size), and on down to atomic scale (where an atom is your size). Micromanipulators will allow actions in these microworlds. Among the things to do down there are to build things (nanotechnology), to explore and probe (biological and physical science), and to interact with microscopic creatures (entertainment).

4.1.6 Global Experience Database and Time Travel. To augment their native memories, all people may in the future come to wear multisensory recording devices, much as people today wear eyeglasses. In addition there will be many other sensors that are simultaneously scanning the real world and recording these data. The data from the sensors at these many locations may be integrated into a global 3-D database, spanning the experience of all of humanity since people began to wear the recorders. It will be possible, using this database, to share the experience of a distant person in real time, or to relive the experience of any person from any time in the past so long as the experience is recorded in the database. This is, in effect, time travel into the recorded past. Time travel into the future is possible to the extent that a simulation can predict what will occur, but there are fundamental theoretical limitations about how well we will ever be able to precisely predict the future. To the same degree that future events can be accurately predicted by simulation, it should be possible to travel into the simulated past and be given the ability to perform actions, and thus to experience what might have been.

4.1.7 Simulating the Real World and Fantasy Worlds. There will be two main kinds of simulated experience: simulations that attempt to accurately mimic or predict what can happen in the real world, and convincing fantasy worlds that dispense with the constraints and physical laws of reality. The accurate simulations will be useful for education, for training, and for exploring the consequences of contemplated actions. The fantasy worlds will entertain, and perhaps delude and addict. Many human participants will be able to simultaneously inhabit these simulated worlds, seeing and interacting with each other. In addition, simulated creatures will also inhabit these worlds. Simulated creatures, also known as autonomous agents or artificial intelligences (AIs), are today quite limited in capability as compared with human beings or even the most primitive animals. At the simplest, a creature is an object in the virtual world that moves about on its own, initiating actions. More sophisticated creatures would include in their behaviors simulations of the abilities to recognize, to remember, and to plan.

4.1.8 Overlays onto the Real World. There will be many databases registered with the real world and able to be superimposed onto it, for example, labels, maps, notes to specific people, diagrams, paths, grafitti, as well as the actions from earlier times recorded in the experience database. It will be a matter of choice which, if any, of these overlays are viewed by each human at any given moment.

4.1.9 Shared Virtual Worlds. Many people will be able to enter simultaneously into these virtual worlds, including real-time transmitted virtual worlds (by virtual travel to the same location), and recorded experience (by traveling to the same place and time in the experience database). For two people simultaneously reliving a particular trace through recorded space and time, it will be a matter of choice whether they see one other. They could equally well relive the same experience independently, or see representations of one another as they observe the action from separate locations. It would also be possible to observe the traces of earlier observers of a given recorded action.

4.1.10 Animals in Virtual Worlds. Humans will be able to masquerade in the real world as animals of any species for which teleoperated robot manikins can be built. To interact effectively with these animals, an understanding of how they communicate is necessary, and this understanding may be facilitated by the very existence of such robots. A particularly effective way to make, for example, a teleoperated cockroach, would be implant remote controls into its neural tissue. This would be a teleoperated but biologically real creature—a zombie cockroach. Zombies would probably be better at fooling animals than the most sophisticated puppet mechanism. However, if cockroaches, frogs, and squirrels can be wired as zombies, then most likely humans can be also.

It will also be possible for higher animals, such as mammals, to enter into virtual worlds in the same way that humans do. If we can reproduce experience with perfect fidelity for humans, then we should be able to do the same for a cat or a snake, with the displays and input devices tailored for their particular senses and physiognomies. Thus, an ape or dog or cat could perceive a surrounding 3-D world, move through it, and interact with it, more or less like they do in the real world. It would be interesting to see if an ape, given the ability to capture and replay representations of its experience for other apes, could make use of this ability for communication.

4.1.11 Direct Neural Connection. Since all perception and action are accomplished through the human sensory and motor nerves, display devices could stimulate the nerve fibers directly, rather than the sensory organs, and in principle achieve the same perception. Likewise, motor nerve impulses could bypass the muscles and directly trigger actuators that manipulate the world.

4.2 Ultimate Limits

Having imagined many grandiose extrapolations to our present capabilities for synthetic experience, we turn to the question of ultimate limits. What aspects and laws of the physical universe and the structure of the human body and brain will constrain what kinds of synthetic experiences are achievable?

4.2.1 Fidelity. For the fidelity of reproduced experience, it seems possible in principle both to cover all human senses simultaneously with display devices and to achieve arbitrarily high fidelity for each sensory channel. For hearing, CD-quality stereophonic sound is already approaching the point of indistinguishability from natural sound. For vision, improvements in video can be posited that would reach the limits of human visual acuity, field of view, color pallette, motion detection, and depth perception. This will not be easy, but neither is it impossible.

Taste has four dimensions (salty, sour, sweet, and bitter) and arbitrary tastes may be synthesized with combinations of these primaries, just as arbitrary colors are synthesized from the red, green, and blue phosphor dots of the television screen. Similarly, smell appears to have seven dimensions (Kandel & Schwartz, 1985), although there is some scientific dispute about this. Thus, smells could probably be synthesized from primary components also.

Displaying to the vestibular system presents more of a problem—motion platforms can tilt and can apply strong but brief accelerations, within the limits of their travel, and imperceptibly drift back to center position. However, displaying a zero-G experience or a 2-min-long 3G Apollo blast-off would seem to require more than a motion platform, perhaps a simulation chamber in high Earth orbit that could be left floating or accelerated at arbitrary rates for long periods.

Displaying with perfect fidelity to the haptic and tactile senses presents such a daunting engineering challenge that it seems nearly impossible. Some sort of whole body exoskeleton would be needed, with integrated arrays of pressure, vibration, and temperature displays covering the entire body surface. However, pieces of this have been done already. The Jacobsen Arm, developed by Steve Jacobsen at the University of Utah, is an exoskeleton with force feedback that covers all the joints of the arm and hand including shoulder, elbow, wrist, thumb, and fingers. A full body force-feedback exoskeleton can therefore be imagined. Tactile arrays that can display texture and vibration to the finger or other surfaces of the skin exist (Linvill, 1973; Rheingold, 1991). Building a flexible, body-covering tactile array would be very difficult, but not impossible.

This completes the list of external sensory organs, leading to the conclusion that experience can potentially be reproduced with a fidelity that is indistinguishable from natural experience. However, a great deal of effort and expense would be required to develop some of the required display devices.

4.2.2 Transmitted Experience. The speed of light limits real-time transmitted experience. In teleoperating a distant robot, there is a speed of light time lag proportional to the remoteness of the robot. This time lag applies both to the transmission of the sensor data to the operator's display and the transmission of the operator's actions to the remote actuators. Humans are sensitive to such time lags, with a lag of 1 sec causing enough confusion to the operator to make many manual tasks impossible. A lag of 100 msec is quite noticeable, and some experiments show sensitivity down to 5 msec. A light-speed circumnavigation of the Earth takes roughly 150 msec, so teleoperation is feasible anywhere on the planet, or in low Earth orbit. However, the lag will begin to be noticeable for operator–robot distances of about 1000 km.

With a 3-sec round trip speed of light delay from the Earth to the Moon, teleoperation is not feasible for dextrous manual tasks that require force-feedback or continuous guidance. However, some things can be accomplished by greatly slowing down the operations performed. To Mars, the minimum round trip time lag is about 10 min, so real-time teleoperation from Earth to Mars is impossible.

4.2.3 Recorded Experience. The fundamental limit on recorded experience is storage capacity. We can estimate the data rate for human experience by using the standard NTSC video data rate for comparison. A hand-held videotape can currently store an hour's worth of visual and auditory experience.

In the future, fidelity and therefore data rate will increase for recorded experience, but at the same time storage density will increase. Let us explore the likely changes in these two parameters. NTSC video nominally transmits an image frame of 640×480 pixels at a rate of 30 frames/sec with the color of each pixel encoded by 8 bits for each of the primary colors red, green, and blue. This is roughly 200 million bits/sec. To be conservative, we will increase this by a factor of 250 to allow for the greater resolution, field of view, and so forth, needed for perfect visual fidelity. Since humans are primarily visual creatures, another factor of 2 should be sufficient to record all the other senses. This gives us a data rate for human experience of 10^{11} bits/sec, which is probably much more than is needed.

Current common storage techniques, such as music compact discs, can store roughly a 10^9 bits/cm^3. However, we can expect storage density to continue to increase until it hits some sort of physical limit. It should ultimately be possible to encode information in the arrangement of matter on the atomic scale (Feynman, 1960; Drexler, 1991). Assuming a nanotechnological storage device that stores 1 bit for every 1000 atoms gives us a storage density of 10^{20} bits/cm^3. At this density, a lifetime (100 years) of human experience can fit into the volume of a large grape (3 cm^3), and the experiences of all of humanity (10 billion people) for 10,000 years would fit into a cubic mile of nanostorage.

Storing a continuous recording of every human's egocentric experience is thus possible in principle. But what about storing data from the many other sensors scattered throughout the world, so as to be able at a later time to review what happened at a location which no one observed at the time of interest? This is possible to a certain degree, but it is impossible to record everything that happened to an arbitrary degree of accuracy. Reality is too complex—even if it were possible to sense the positions of all the atoms in a scene from nanosecond to nanosecond, this would overwhelm even the nanostorage postulated above. It does not seem possible to record everything that one might want to later experience. Recording will necessarily be selective.

4.2.4 Simulated Experience. The fundamental limit on the accuracy of simulations of reality is the inherent complexity of reality itself. Also, our current understanding is that reality is fundamentally unpredictable at the quantum level, and therefore at all higher levels. We can expect simulations that attempt to predict and model reality to achieve increasing fidelity within restricted domains (such as vehicle simulators, weather prediction, or orbital mechanics), but there can be no expectation of being able to predict the *exact* behavior of macroscopic quantities of matter under arbitrary conditions. There are several reasons for this. Chaotic systems existing in nature are sensitive to tiny perturbations and therefore defy prediction. The uncertainty principle bars accurate knowledge of the parameters needed for prediction of behavior on the atomic scale. Certain quantum events appear to be truly unpredictable. Even without all the foregoing reasons, the sheer number of atoms involved makes such an atom-by-atom simulation infeasible.

It would seem that it will never be possible to predict the exact behavior of an ensemble of atoms because of the reasons of chaos, uncertainty, and quantum unpredictability. And even without those objections, it would also seem that it would take a simulation computer enormously larger in physical size than the object of the simulation, and that the simulation would proceed much slower than the atomic interactions being simulated. But to be cautious, we might phrase it as a question: Does it always take more atoms and more time to simulate (and predict the outcome of) a physical process than are involved in the process itself?

This limit on simulation accuracy is a fundamental limit, not just a temporary obstacle that can be overcome with effort and time. How can simulated experience deal with this? Several techniques have proved useful so far. Selective fidelity looks at the limited domain being simulated and the task to be done to select which aspects of experience to simulate accurately. Interactive steering of the display computation dynamically allocates computational resources to where they will be most effective, as in, for example, multiple levels of detail for objects in flight simulators that are chosen according to the user's current location.

For fantasy worlds that do not attempt to accurately model the real world, it is hard to see a fundamental limit. With the same display apparatus that can reproduce transmitted experience with perfect fidelity, the display should not impose any limit. Extrapolating into the future, we can expect enormous increases in computational power available for running the simulations underlying these fantasy worlds. Perhaps the limitation will simply be a noticeable difference between the way things happen in simulated worlds and the way things happen in reality, which cannot be simulated in detail.

5 Conclusions

The taxonomy presented in this paper, with its nine dimensions, offers a way to classify devices that use technology to transmit, filter, record, or simulate experience. The taxonomy also helps to understand the relationships among existing synthetic experience devices, and to suggest as-yet untried possibilities.

It appears that the new capabilities offered by the various kinds of synthetic experience can be developed rather far before we hit fundamental limits. It seems plausible that in the near future, we will be able to transmit experience from a distance with good fidelity and coverage of most senses, see the invisible through sensors linked to HMDs, travel visually to distant places as easily as we make telephone calls, record 3-D scenes and actions for later replay, and enter simulated 3-D worlds. These simulated worlds may be either fantasy worlds or may try to accurately model some aspect of the real world.

The speed of light limits the distance that experience may be transmitted without perceptible time lags. The accuracy with which the real world can be simulated is limited by several factors that make it unlikely and probably impossible to predict by simulation the *exact* behavior of many aspects of the real world. The recording of experience is limited by storage capacity. If storage of information can be realized in the arrangement of matter at the atomic scale, continuous recording of human experience may be feasible.

Acknowledgments

I would like to thank my colleagues in the UNC Computer Science Department for their help and support, particularly the Head-Mounted Display team (Fred Brooks, Henry Fuchs, Gary Bishop, Rich Holloway, Jim Chung, Drew Davidson, Erik Erikson, Mark Mine, Jannick Rolland, and Doug Holmgren), the Nanomanipulator team (Stan Williams and Rick Snyder of UCLA, Bill Wright, Vern Chi, and Russ Taylor), and the Pixel-Planes team. Fred Brooks, Russ Taylor, and Ivan Sutherland suggested modifications to earlier versions of the classification system presented in this paper. The anonymous reviewers also made many helpful suggestions and corrections. I would also like to thank my friends and colleagues from the Banff Centre for the Arts, NASA Ames Research Center, The Learning Company, and Atari, where my ideas about computer-simulated worlds evolved. I would like to thank Julius Smith, Scott Kim, Fred Lakin, Ken Harrenstien, and Ken Shoemake for many years of technical discussions, critiques, and support. This research was supported by the following grants: DARPA DAEA 18-90-C-0044, NSF Cooperative Agreement ASC-8920219, ONR N00014-86-K-0680, and NIH 5-R24-RR-02170.

References

Blanchard, C., Burgess, S., Harvill, Y., Lanier, J., Lasko, A., Oberman, M., & Teitel, M. (1990). Reality built for two: A virtual reality tool. *Proceedings of the 1990 Workshop on Interactive 3-D Graphics,* 35–36.

Bright, M. (1984). *Animal language,* 55–108. London: British Broadcasting Corporation.

Caudell, T. P., & Mizell, D. W. (1992). Augmented Reality: An application of heads-up display technology to manual manufacturing processes. HICSS conference, January 1992.

Clarke, A. C. (1962). *Profiles of the future.* New York: Holt Reinhart & Winston.

Drexler, K. E. (1991). *Unbounding the future.* New York: William Morrow.

Feynman, R. P. (1960). There's plenty of room at the bottom. Reprinted in *Journal of Micromechanical Systems 1*(1), 1992.

Gruss, A., Tada, S., & Kanade, T. (1992). A VLSI smart sensor for fast range imaging. *Proceedings of IROS '92,* July 7–10.

Kandel, E. R., & Schwartz, J. H. (1985). *Principles of neural science* (pp. 418–419). New York: Elsevier Science Publishing.

Linvill, J. G. (1973). *Research and development of tactile facsimile reading aid for the blind (the Opticon).* Report to U.S. Department of Health, Education, and Welfare, Stanford Electronics Laboratory, Stanford University.

Naimark, M. (1991). Elements of realspace imaging: A proposed taxonomy. *Proceedings of SPIE: Electronic Imaging,* 1457.

Naimark, M. (1992). *Elements of realspace imaging.* Tech. Rep., Apple Multimedia Lab, Apple Computer, Cupertino, CA, 44 pages.

Rheingold, H. (1991). *Virtual reality.* New York: Summit.

Robinett, W. (1991a). Electronic expansion of human perception. *Whole Earth Review,* Fall issue, 16–21.

Robinett, W. (1991b). Technological augmentation of memory, perception, and imagination. *Virtual Seminar on the Bioapparatus,* 17. Banff, Alberta, Canada: The Banff Centre for the Arts.

Robinett, W., Taylor, R., Chi, V., Wright, W. V., Brooks, F. P. Jr., Williams, R. S., & Snyder, E. J. (1992). *The nanomanipulator project: An atomic scale teleoperator,* 1992 SIGGRAPH course notes on "Implementation of Immersive Virtual Worlds."

Weiss, R. (1989). New dancer in the hive. *Science News 136*(18), 282–283.

Zeltzer, D. (1992). Autonomy, interaction, and presence, *Presence 1*(1), 127–132.