David Rokeby

The Construction of Experience : Interface as Content

This article appears in the book:
"Digital Illusion: Entertaining the Future with High Technology," Clark Dodsworth, Jr., Contributing Editor
© 1998 by the ACM Press, a division of the Association for Computing Machinery, Inc. (ACM)
published by Addison-Wesley Publishing Company.

INTRODUCTION

   I'm an interactive artist; I construct experiences. Since the early 80's I've been exhibiting my installations in galleries, trade shows, science museums, and public and private spaces. These exhibitions serve as a public research laboratory where my ideas about interaction and experience are tested, affirmed, or shot down. This is a condensation of the results of my free-form research.

   Entertainment has traditionally involved heavily coded communication. It has predominantly been delivered through words, sounds, symbols and gestures which stimulate the imagination to render an experience. The visual arts and theatre at various times in history, and film and television in the past century use the direct visual experience of images as a way to make the experience more immediate… to make the audience feel more "there." But these experiences remain things that happen to you. Interactivity's promise is that the experience of culture can be something you do rather than something you are given. This complicates our conventional ideas about "content" in the context of this new medium.

INTERFACES ARE CONTENT

   Everyone is talking about content in interactive media these days. Independent artists and the entertainment industry alike now see that these new technologies are relatively flat without significant content. But the rush to stuff content into interactive media has drawn our attention away from the profound and subtle ways that the interface itself, by defining how we perceive and navigate content, shapes our experience of that content. If culture, in the context of interactive media, becomes something we "do," it's the interface that defines how we do it and how the "doing" feels. Word processors change the experience of writing, regardless of the content; they affect the manner in which that content is expressed. Hypermedia provide multiple trajectories through content, but the nature of the links, branches and interconnections influences our path, and inevitably changes our sense of the content. Active agents, either in our software or on the net, guide us through the information jungle; they're sorting demons, deciding what's relevant and irrelevant, providing us with interpretation and point-of-view. Marshall McLuhan's phrase "the medium is the message" became a tired cliché long before our media became flexible and intelligent enough to live up to the epithet. Like most cliches, it carries plenty of truth, and needs a full re-examination in the context of emerging active and interactive technologies.

THE TRICKS OF THE TRADE

   The creation of interactive interfaces carries a social responsibility. I've come to this conclusion from my experience creating and exhibiting interactive systems. At first glance, it may seem that I'm stretching the point here. It's really just entertainment, right? Indeed, as an artist, it's my traditional right to use every trick in the book to create a magical experience. Fantasy and illusion are key elements of most effective culture, from high-brow theatre to video games. Hollywood has always relied on sets, stunt people, and special effects to get its stories across. Computer game developers are the newest masters of illusion.

   One of the clearest examples I can recall is an early videodisc-based video game in which users got the impression that they were flying at great speed over a terrain. The videodisc was made up of video clips which linked together in a branching and merging structure. The image I saw on the screen was the middle portion of the full video frame. If I turned to the left during a linear video segment, the section of the frame that I saw instantly panned in that direction, giving an immediate sense of responsiveness, but I was, in fact, still travelling along the same restricted path. The illusion that I had the freedom to roam the entire terrain was maintained for a surprisingly long time, partly because I was moving at a high 'virtual'

speed without time to reflect on the degree to which my actions were being reflected. That technique was a brilliant and effective way to get around the inherent limitations of videodisc as a real-time interactive medium. Whether you really had freedom to wander the terrain was beside the point, because the game was engaging and exciting.

The line between entertainment and everything else is getting very vague these days (infotainment, edutainment). The Web represents a convergence of the video game industry and commercial transaction systems, and this leads to a potential problem; illusion translated into the commercial world becomes deception. The tricks of today's artists and hackers are the commercial tools of tomorrow. Perhaps more significantly, with the explosive growth of the internet, these sleights-of-hand are becoming incorporated into communications systems, and by implication, into oursocial fabric. Whether we intend it or not, we're redesigning the ways that we experience the world and each other.

VIRTUAL SPILL
There are two levels of leakage here. On one level, there is the effortless migration of code and hardware from the entertainment world to the "serious" worlds of commerce, justice, and communication. At the second level, artificial experiences subtly change the way we feel, perceive, interpret, and even describe our "real" experiences.

The most graphic and extreme example of virtual spill into the real is probably VR-sickness, an after-effect of Virtual Reality. My experience was that I would suddenly lose my orientation in space at apparently random moments for about 24 hours after my virtual immersion. I felt as though I were off the floor, and at an unexpected angle. As far as I can tell, the explanation was that, when I was immersed, I'd desensitized my response to the balancing mechanisms in my inner ears in order to sustain the illusion of motion in a purely visually defined 3D space. Once I was desensitized, I was free to accept the illusion of space that the VR system provided. But on returning to "real" space, my inner ears didn't immediately resume their job. I was taking my sense of orientation in space entirely from visual cues. One attack may have been stimulated by a design of sharply angled lines painted on a wall. My visual system seems to have interpreted this cue as vertical, and abruptly changed its mind about my body's orientation, while my ears were certain that I was standing quite straight, bringing on a wave of nausea.

I've also experienced after-effects from spending extended periods interacting with my most- exhibited interactive installation, Very Nervous System. In this work, I use video cameras, an artificial perception system, computer, and synthesizer to create a space in which body movements are translated into sound or music in real-time. An hour of the continuous, direct feedback in this system strongly reinforces a sense of connection with the surrounding environment. Walking down the street afterwards, I feel connected to all things. The sound of a passing car splashing through a puddle seems to be directly related to my movements. I feel implicated in every action around me. On the other hand, if I put on a CD, I quickly feel cheated that the music does not change with my actions.

When I first got a Macintosh computer and spent endless days and nights playing with MacPaint, one of the things that amazed me most was the lasso tool, which allows you to select a part of the image and drag it across the screen to another location. The most intriguing thing was the automatic clipping of the background behind the dragged selection. Walking down the street after an extended MacPaint session, I would find myself marvelling as backgrounds disappeared behind trees, acutely aware of what was momentarily hidden from view.

Interfaces leave imprints on our perceptual systems which we carry out into the world. The more time we spend using an interface, the stronger this effect gets. These effects can be beneficial or detrimental. Dr. Isaac Szpindel at the Jewish General Hospital in Montréal is experimenting with the use of "Very Nervous System" as a therapy for Parkinson's disease. People suffering from this disease tend to lose their ability to will their own movement, but remain capable of responding quickly in emergencies. While the results are still preliminary, it appears that regular interactions with Very Nervous System can help to re-engage Parkinson's suffers' ability to motivate their own movement in their normal day-to-day lives.

CONCEPTUAL SPILL

Exposure to technologies also change the ways that we think and talk about our experiences. We use terms borrowed from computers when describing our own mental and social processes. We "access" our memories, we "interface" with each other, we "erase" thoughts, we "input" and "output." In a chillingly insightful comment on the way technologies and ideas interact, Alan Turing, one of the great computer pioneers wrote: "I believe that at the end of the [twentieth] century, the use of words and general educated opinion will be altered so much that one will be able to speak of machines thinking without expecting to be contradicted."[1]

This statement is often taken to mean that Turing believed that machines would be able to think by the turn of the century. In fact, he is saying that our ideas of what thinking is and what computers can do will converge to the point that we cannot express or grasp the difference. This sort of convergence may also soon take place in the realm of experience; we may lose our ability to differentiate between raw and simulated experience.

In 1983 I was invited at the last minute to exhibit my interactive sound installation in an exhibition called "Digicon '83" in Vancouver. This was to be my first public show, and I was very excited, but there was a tremendous amount of work to be done. I worked between 18 and 20 hours a day refining an interactive interface from a barely implemented concept to an actual experiential installation. I spent no time with friends and didn't get out at all. I got the piece done and was extremely pleased with the results. After setting up my installation in Vancouver, I was astonished by the fact that it did not seem to respond properly to other people, and sometimes didn't notice people at all. I didn't really understand the problem until I saw videotape of myself moving in the installation. I was moving in a completely unusual and unnatural way, full of jerky tense motions which I found both humorous and distressing. In my isolation, rather than developing an interface that understood movement, I'd evolved with the interface, developing a way of moving that the interface understood as I developed the interface itself. I'd experienced a physiological version of the very convergence that Turing described.

While we may lose our ability to understand and articulate the differences, we will still have some intuitive sense of them. But many of the differences between virtuality and reality will be subtle and easy to discount, and intuition often loses in the face of hard logic; we may find it as easy to ignore our intuitions as to ignore our inner-ears while immersed in VR. I believe there are important reasons, beyond simply romantic nostalgia, to nurture awareness of the distinctions between the real and the virtual.

THE EXPERIENCE OF BEING

By defining a way of sensing and a way of acting in an interactive system, the interface defines the "experience of being" for that system. Through their design of the interface, the creators have in large part defined the user's "quality of life" while they are interacting with the system. Unfortunately, the design parameters for quality of life are pretty undefined. There seems to be no agreement on what makes for a high "quality of life." I suspect it's dependent on a whole range of parameters that we rarely pay attention to.

In order to better understand what those parameters are, we need to look at how our experience of the real world is constructed. In other words, what is our user interface for reality? or: What is the nature of our relationship with the world? I don't intend?nor am I qualified?to plumb the depths of philosophical thinking on this subject. There is a branch of philosophy dedicated to these questions called "phenomenology" for those who want to explore this in greater depth.

THE BANDWIDTH OF "REALITY"

Our "organic" interface is extraordinarily complex and massively parallel. Our sensing system involves an enormous number of simultaneously active sensors, and we act on the world through an even larger number of individual points of physical contact. In contrast, our artificial interfaces are remarkably narrow and serial even in the multimedia density of sound and moving image. These interfaces are also unbalanced in terms of input and output. At the computer screen, we receive many thousands of pixels at least 60 times a second from our monitors, while sending a few bytes of mouse position or keyboard activity back to the system. We appear to most of our interactive systems as a meagre dribble of extremely restricted data. Even in immersive VR systems, we're commonly represented as a head

orientation and a simple hand shape. We may imagine ourselves immersed in the Virtual Reality, but the Virtual Reality is not, from its point-of-view, enveloping us.

It's not simply a question of lack of senses such as touch and smell. It's also a question of the actual number of contact points through which an interaction passes. Our nervous system, senses and perceptual systems integrate an enormous number of separate inputs in order to construct our sense of being. The "bandwidth" of real experience is almost unimaginable. In order to fit interaction into the available bandwidth of our computers and communications systems, we must decide what narrow aspects of the user's presence and actions will be involved. It's an extreme form of compression, and it's "lossy."

BOTTOM-UP VERSUS TOP-DOWN

Through our human interface, we access a pool of content of unimaginable complexity. This content exists at many levels. There is raw sensual content. There are people and things and their complex behaviours and interactions. There are conjunctions of actors and actions that play out in an apparent causality. There are stories, symbols, words and ideas. While our attention is often focussed at one or another of these levels, our sense of being is constructed by input at all of these levels simultaneously.

The whole system is built from the ground up. Subatomic particles interact to produce atoms and molecules. Atoms and molecules interact to produce organic and inorganic matter. Matter gathers into things with higher order behaviours like mountains, rivers, plants and humans. These things interact in the whole complex process of life. Ideas, words and concepts are things that we use to describe these processes. They're inexact generalizations and simplifications that are necessary for our sanity.

Artificial interfaces may access a pool of information as large as the internet, but the internet is tiny compared to reality. And that pool of content generally starts at the level of words and ideas. A digital image is similarly abstracted. It's not self-generated from the interactions among its constituent pixels.

Artificial experiences are built up as a sort of collage of representations of things torn out of context. In the virtual realm, context is purely a matter of the taste of the creator. One decides arbitrarily to put these things together. In "reality," the context is not just the ground against which you see something; it represents the set of conditions which makes the presence of the thing possible. The difference is immense, and the more interactive, immersive, or convincing an artificial environment is, the more careful we must be.

Real experience has a fundamental integrity that virtual experience does not. This aspect of virtuality can be a great advantage because it allows you to break the "rules" of reality. Escaping reality is liberating when one spends the greater part of one's time in reality. But this lack of fundamental integrity is potentially quite unsettling to anyone spending most of their time in virtual spaces.

A HARDENED PERCEPTUAL EXOSKELETON

The input from our senses generally reaches our awareness only after passing through the powerful filters of our perceptual systems, but we can also open ourselves to raw sensuality. There is something profoundly important about the fact that the base of our human/reality interface is raw and uncoded. We can, to some degree, bypass our own perceptual filtering.

I had an experience in art school that brought this home in a very direct way. One of my professors told us one day that we would be looking out a window for the whole three-hour class. I was incensed. I'd been willing to go along with most of the unusual activities these classes had entailed, but I felt this was going too far. I stood at my assigned window and glared out through the pane. I saw cars, two buildings, a person on the street. Another person, another car. This was stupid! For fifteen minutes I fumed, and muttered to myself. Then I started to notice things. The flow of traffic down the street was like a river, each car seemingly drawn along by the next, connected. The blinds in each of the windows of the facing building were each a slightly different colour. The shadow of a maple tree in the wind shifted shape like some giant amoeba. For the remaining hours of the class I was electrified by the scene outside. After fifteen minutes, the "names" had started separating from the objects.

It seems that we stop seeing, hearing, smelling as soon as we have positively identified something. At that point, we may as well replace the word for the object. Since identification usually happens quickly, we

spent most of our time not really sensing our environment, living in a world of pre-digested and abstracted memories.

This explains our attraction to optical illusions and mind-altering experiences (chemically-induced or not). Those moments of confusion, where identification and resolution aren't immediate, give us a flash of the raw experience of being. These moments of confusion are also the fulcra of paradigm shifts. It's only when our conventional way of dealing with things breaks down that we can adopt another model, another way of imagining and experiencing a scenario.

The adrenaline rush of a high-speed video game has something in common with this experience of filter-breakdown; the barrage of images and the need to act quickly test the limits of our perceptual and responsive systems. But these systems have then added themselves to our internal filters and they aren't subject to this same sort of breakdown. Paradigm shifts in the interface can only happen through the software and hardware development cycle, which is burdened with economic considerations and intense industry competition. The interface becomes a hardened and brittle perceptual exoskeleton which we can't easily question or redefine. This becomes increasingly problematic as the interface becomes more "transparent" and "intuitive." At those difficult and confusing moments when our way of viewing the world needs to change, we may not know to examine the interface as a potential contributor to the problem. For this reason, I believe we need to develop an interactive literacy. We need to learn, then teach others to critique and understand the influences of our interfaces as we use them.

PUNCH AND SCREAM

Our interface with reality is not only multi-sensory or multimedia but also "multi-modal." We can talk, scream, gesture or punch. We can interpret, analyze or simply enjoy the raw sensation. It's only a multi-modal approach with multiple simultaneous levels of meaning and communication that can properly express that complex experience of reality.

In 1988, I was invited to exhibit "Very Nervous System" at the Siggraph Art Show. "Interactivity" was just emerging as a buzzword and there was a lot of scepticism on the floor. Many attendees entered my installation to "test" it using what I've come to call the "First Test of Interactivity." The test involves determining whether the system will consistently respond identically to identical movements. (Note that an intelligent agent will probably fail this test.) They would enter the space, let the sounds created by their entrance fade to silence, and then make a gesture. The gesture was an experiment, a question to the space; "What sound will you make?". The resulting sound was noted. Second and third gestures were made with the same motivation, and the same sound was produced. After the third repetition, the interactor decided that the system was indeed interactive, at which point they changed the way they held their body and made a gesture to the space, a sort of command: "Make that sound." The command gesture was significantly different from the early "questioning" gestures particularly in terms of dynamics, and so the system responded with a different sound. I observed a couple of people going through this cycle several times before leaving in confusion. Their body had betrayed their motivation.

Body movement can be read on two different levels. There is the semantic content of the gesture, in which the movement is interpreted symbolically (the "OK" sign or the raised middle finger), and there is the raw visual experience of the gesture, to which my system was responding. The questioning and commanding gestures were semantically similar but quite different in terms of physical dynamics. More practical interactive interfaces might filter out the involuntary dynamics of the gesture, treating them as unwanted noise, and focus on the semantic content. In interpersonal communications, we're always simultaneously interpreting gestures on many levels. This combination is the basis for richer communication. For this breadth and quality of communication to be carried through interfaces, the designers must be aware of the importance of these multiple modes, and then must be able to actually create the code and hardware to support them.

A multi-modal interface would be particularly important in engendering trust and intimacy through communications systems. Sweat, smell, nervous gestures, cold or warm hands, tone of voice, exact direction of gaze are all elements by which we gauge subtle interpersonal conditions like honesty, nervousness. I'm not advocating interfaces involving every possible sense; the literal replication of the whole nervous and sensing systems would be cumbersome and overly literal. I'm just pointing out the

many complex levels that exist in real flesh-to-flesh communication. We need effective ways to accommodate simultaneous layers of communication if our telecommunication is to be satisfying and richly successful.

HUMAN INTERFACES AS BELIEF SYSTEMS

So there are very large differences between the human and artificial interfaces. Quite often, the simplified, symbolic nature of the artificial experience is a useful characteristic. This is particularly true for situations that involve abstractions like numbers, words or ideas. The interface in this case clearly suits the material and can make those abstractions more accessible through simulation and visualization.

But we're spending more and more time amongst our simulations, and we're in danger of losing sight of the fact that our models and ideas of "reality" are drastically simplified representations. If we do lose this awareness, then our experience of being will be significantly diminished. Simulations offer us formerly unimaginable experiences, but the foundations of these simulations are built up from a relatively narrow set of assumptions about the structure and parameters of experience. And the built-in exigencies of product development mean that this narrow set of assumptions and ideas quickly becomes a standard, and soon after, crystallizes into silicon for performance gains. Once there is an inexpensive chip available, these assumptions have become practically unassailable for a considerable length of time.

In an odd way, this parallels medieval Christianity. During the middle ages, the church sanctioned a certain set of ideas about the world. This system of beliefs became the standard "browser" for viewing reality. Many of the assumptions built into that system were clearly absurd from our contemporary point of view, but they had a grip on the imagination of the whole Christian world, to the point that brilliant philosophers went through ridiculous contortions to justify officially-sanctioned ideas that seem to us ridiculous. The interface designers of this era were the monks, bishops and popes.

Our user interfaces are also a kind of belief system, carrying and reinforcing our assumptions about the way things are. It's for this reason that we must increase our awareness of the ways that the interface carries these beliefs as hidden content. It may be hard to conceive of the standard GUI as a belief system, but the "holy wars" between Macintosh and Windows users on the internet indicates an almost religious passion about interface. It's also useful to realize that effective interfaces are usually intuitive precisely because they tap into existing stereotypes for their metaphors. An interface designed for racists might tap into racist stereotypes as a source for icons and metaphors that would be immediately understood across the user-base. A metaphoric interface borrows cliches from the culture but then reflects them back and reinforces them.

BEYOND LITERAL SIMULATION

I've argued that virtual experiences don't do justice to the richness of the human experience. I'm not suggesting that the richness of experience cannot be increased through interactive technologies or that the best interface would be one that exactly replicated the full experience of reality. In fact, designers of virtual experiences are often so literal in their attempts to simulate reality that they stifle some of the most exciting potentials that these new media offer.

There is an artist named Tamas Waliczky who has been working with the idea of alternate systems of 3-D representation. The conventional binocular, perspectival model that is currently being standardized in software and hardware is useful for normal representations of 3-D objects and space. The fact that it has reached the level of silicon represents the kind of crystallization I mentioned above. Waliczky sees much broader possibilities for the representation of space than this limited Renaissance model. He has been creating alternate perspectival systems, writing code to render other experiences of space. In one of his works, he renders a world from the self-centred point-of-view of a young child. For another, he has created a program that renders inverted perspective, in which things get larger as they get further away from you, and vice versa. This is a real mind-bender to see, a rich exploration of the potentials of virtual media to go beyond the restrictions of reality, and indeed of our own imaginations.

I'm suggesting that there is a sort of middle-of-the-road virtuality, that does justice to neither our rich experience of reality, nor the richness of possible virtual experience. This 'MOR' virtuality diminishes

experience in several dimensions without enhancing it in others. If virtual experiences are to add to the dimensions of experience we must avoid this mushy middle ground when imagining and designing them.

DESIGNING THE EXPERIENCE OF BEING

How does an interface form the experience of being? How do the decisions of the designer and programmer shape the experience of the user? I'll try to examine these by looking at several general characteristics or "qualities" of the interface.

THE DISTORTING MIRROR

An interface inherently constructs a representation of the user. To a computer with a simple graphic user interface, the user is a stream of mouse clicks and key taps. Advanced interfaces involving intelligent agents compile much more detailed representations of the user by interpreting this stream of input and attempting to determine the intentions of these activities so that the interface can help the user be more efficient. How the user is represented internally by the system defines what the user can be and do within the system. Does the system allow the user to be ambiguous? Can they express or act upon several things at once?

Interactive systems invariably involve feedback loops. The limited representation of the user is inevitably reflected back to the user, modifying their own sense of self within the simulation. The interface becomes a distorting mirror, like those fun-house mirrors which make you look fat, skinny or a bizarre combination of the two. A standard GUI interface is a mirror that reflects back a severely misshapen human being with large hands, huge forefinger, one immense eye and moderate sized ears. The rest of the body is simply the location of backaches, neck strain, and repetitive stress injuries. It's generally agreed that the representation of women or visible minorities in magazine and television advertisements affects their self-image. If we accept this, then we must also accept that interface-brokered representations can exert a similar, though more intimate, effect on the reflected computer user.

THE CONSTRUCTION OF SUBJECTIVITY

We are who we are, with a unique character and personal idiosyncracies, largely because of our individual subjective viewpoint. That viewpoint is formed by our windows out into the world (our senses), in conjunction with our memories and experiences. An interactive interface is a standardized extension that shapes and modifies the user's subjective point of view. By presenting information in a specific manner or medium, the interface designer defines the way this interface shapes this point-of-view.

There is a paradox in the manner that interactive systems affect our subjectivity. The non-interactive system can be seen as stubborn in refusing to reflect the presence and actions of the spectator, or, it can be seen as giving the spectator complete freedom of reflection and interpretation by not intervening in the process. An interactive system can be seen as giving the user the power to affect the course of the system, or as interfering in the interactor's subjective process of exploration. An extreme example would be an interactive system which detected whether the user was male or female and presented different content or choices to members of each sex. The system would be closing off parts of the system to each person due to their gender. Such an interactive system displaces some of the user's freedom to explore the content. Any interactive interface implicitly defines the "permissible" paths of exploration for each user.

This irony gets increasingly pronounced as the technology of interaction becomes more sophisticated. In the introduction to his book, Artificial Reality II, Myron Krueger invites us to "Imagine that the computer could completely control your perception and monitor your response to that perception. Then it could make any possible experience available to you." [2] Florian Roetzer responds that a system that gives you this "freedom" of experience must necessarily be a system of infinite surveillance.[3] When a system monitors its users to this extent, it has effectively taken control of their subjectivity, depriving them of their idiosyncratic identity and replacing it with a highly focussed perspective that is entirely mediated by the system. Subjectivity has been replaced by a synthetic subjective viewpoint. The fact that the system responds to the interactor does not guarantee in any way that the system is responsible to the user; the interactor can fairly easily be pushed beyond reflection to the edge of instinct, capable only of visceral response to the system's stimuli, mirroring the systems actions rather than being mirrored by the system.

THE INTERFACE AS A LANDSCAPE

Interactive interfaces can explicitly define "permissible" paths of exploration for each user, but in most cases, it's more subtle than that. It's usually not so much a matter of permission as of paths of least resistance. An interface makes certain actions or operations easier, more intuitive, or more accessible. By privileging some activities, it makes the unassisted operations more difficult and therefore less likely to be used. A feature that requires seven layers of dialogue boxes is less likely to be used than one which requires a single keystroke. The interface defines a sort of landscape, creating valleys into which users tend to gather, like rainwater falling on a watershed. Other areas are separated by forbidding mountain ranges, and are much less travelled. A good interface designer optimizes the operations that will be most often used. This practise carries the hidden assumption that the designer knows how the interface will actually be used. It also tends to encourage operational cliches; things that are neat, easy to do, and thus get overused. Software assistants add another layer to this landscape. Like a Tibetan sherpa guiding you up Mount Everest, intelligent assistants make it easier to traverse the more forbidding parts of the landscape, but they themselves create a second landscape. A guide selects and interprets, and may just as easily hide possibilities from you as present them.

DROWNING IN OUR OWN CLICHES
My early interactive sound installations were programmed in 6502 assembly language (the 6502 was the processor in the Apple][). I developed interactions in those days by setting up a simple interactive algorithm, testing the experience for a while, then modifying the code to implement the resulting new ideas. After the programming I'd have to do some debugging. Finally, after up to several hours of work, I could actually step in and experience the alteration. As a feedback loop this process was severely flawed. By the time I'd implemented the idea, I'd often lost track of the idea that sparked the modification. I decided to make the development process as interactive as the experience itself, so I wrote a simple language with which I could modify the behaviours in real-time. I took the basic structures and processes that I'd been coding in assembly language and turned them into standard objects and instructions. This language allowed me to create works in hours that would have taken months to realize in assembly code. It also allowed me to build more complex interactions from these standard building blocks, like any higher-level language. However, it took me a few months to realize that the language was having another effect on my installations. They were becoming less interesting; the building blocks of interaction that made up the language had become cliches.

Assembly code itself contains very little in the way of abstraction. I would take my idea and build it, as it were, atom by atom. It presents a relatively level playing field. My higher-level language was more of a terrain, with peaks and valleys. Once I was placed somewhere on that landscape, there were pathways that were easy and pathways that were difficult. My decisions about how and what to implement were inevitably influenced by this terrain. A landscape gives you a fine view in some directions and obscures others.

Everything that builds on abstractions (languages, perception, and user-interfaces) creates a biased terrain, even as it makes certain previously impossible things possible. Structural differences between languages like Chinese and English subtly cause native speakers to view the world differently. But whereas a spoken language has evolved over centuries and has had millions of unique co-designers from all walks of life, a user-interface or computer language has usually been designed by a small team of people with a lot in common. And they were probably in a big hurry.

SOFTWARE PUBLISHING AS BROADCASTING
When the Apple Macintosh first came onto the market, MacPaint™ sent a shock-wave through the creative community. For the first year, MacPaint-produced posters were everywhere, an explosion of the possibility for self-expression. But while the MacPaint medium reflected the user's expressive gestures, it also refracted them through its own idiosyncratic prism. After a while, posters began to blend into an urban wallpaper of MacPaint textures and MacPaint patterns. The similarities overpowered the differences. Since then, graphics programs for computers have become more transparent, flexible, and commonplace, but the initial creative fervour that MacPaint ignited has abated. The restrictions that made MacPaint easy to use were also the characteristics that ultimately limited its usefulness as a medium for personal expression.

Television, radio and print broadcasting are portrayed as the bad boys from which interactivity rescues us. Interaction allows us to access a wider range of information, not just what the networks choose to broadcast. However, interactive systems do their own kinds of broadcasting; transmitting processes, modes of perception, action, and being. When you define how people access and experience content, you have a more abstract form of control over their information intake. It doesn't matter that every piece of information in the world is on the internet, if the browsers and search engines, through biases in their design, make it unlikely that certain information will be found. It is not difficult to imagine an internet search engine provider selling search priority points; You pay your money, and your company's web pages would automatically get an extra 10% rating on each query in which they come up, putting them closer to the top of the list of results, and so more likely to be accessed.

I've no desire to demonize interactive technologies. But we need to remind ourselves of the ways they subtly shape our experience, particularly in the face of the wild utopian rhetoric that currently surrounds interactiivity. Yes, interactive media can empower and enfranchise. But they simultaneously create new kinds of constraints on abstract and psychological levels, constraints that are more difficult to understand and critique than the familiar biases of the press and broadcast media. Information itself does not create meaning; meaning is created by context and flow, selection and grouping. By guiding us through jungles of content, interfaces are partially responsible for the meanings we discover through them.

FREEDOM VERSUS CONTROL

In the early days of "Very Nervous System" I tried to reflect the actions of the user in as many parameters of the system's behaviour as possible. I worked out ways to map velocity, gestural quality, acceleration, dynamics, and direction onto as many parameters of sound synthesis as I could. What I found was that people simply got lost. Every movement they made affected several aspects of the sound simultaneously, in different ways. Ironically, the system was interactive on so many levels that the interaction became indigestible. People's most common response was to decide that the sounds from the system were not interactive at all, but were being played back on a cassette deck.

I found that as I reduced the number of dimensions of interaction, the user's sense of empowerment grew. This struck me as problematic. I had, at the time, very idealistic notions about what interaction meant (and how it would change the world). In retrospect, the problem seems to have been a linguistic one: people were unfamiliar with the language of interaction that I gave them. Simplifying the language of interaction by reducing its variables let people recognize their impact on the system immediately. With repeated exposure, the user could handle and appreciate more nuanced levels of interaction. In time they could appreciate the flexible, expressive power I'd been trying to offer in the first place.

This is a comforting notion, but it only works if the interactive system stays the same long enough for users to become expert. At the current rate of technology development, such familiarity may never have a chance to develop. As perpetual new users, we may be drawn inexorably toward simplistic systems, trading real power for an ever-evolving glimpse of some never-to-be-achieved potential.

THROUGH THE FEEDBACK LOOP

Interactive systems inherently involve feedback. The system responds to your actions, and you respond based on its responses and your desires. In "Very Nervous System," I constructed tight real-time feedback loops with complex behaviours which illustrated several interesting characteristics of interactive feedback. The responsive character of "Very Nervous System" is built up of little virtual instrumentalists, each of which improvises according to its personal style based on what it "sees" through the camera. Some of those virtual players are drummers, who respond to movement with rhythmic patterns. A rhythmic pattern doesn't necessarily have anything to do with an on-camera person's rhythmic motion, it's merely that virtual player's way of responding. People often involuntarily fall into sync with one of those rhythms, then exclaim that the system is so "intelligent" that it synchronized to their movement!

This illustrates an interesting side effect of real-time interactive feedback loops. An action provokes a response which immediately provokes a shift in action which likewise immediately changes the system's response, ad infinitum. The issue of who is controlling whom becomes blurred. The intelligence of the human interactors spreads around the whole loop, often coming back in ways they don't recognize, causing them to attribute intelligence to the system.

CONSCIOUSNESS LAGS BEHIND THE BODY

Another reason for the confusion between what we as interactors do and what the system does, is that our consciousness seems to trail our actions by up to one tenth of a second. It takes that long for us to be fully aware of what we're doing. I once programmed "Very Nervous System" to respond very clearly as soon as it saw the slightest movement. In every instance, the system responded before I realized that I'd started moving. In fact, the system seemed to respond at the moment that I decided I would move. This delay in consciousness makes it possible for systems with high sampling rates and response speeds to slip under the user's consciousness. At this point, the system and its responses are experienced in the same way that we experience our own body. The interactive system becomes integrated into our proprioceptive system?the same internal sensing system that defines our sense of being in our body, and establishes the relative position of our arms and legs to our "point of consciousness."

This phenomenon, like all the others I describe in this text, cuts both ways. Part of the desire that drove me to produce "Very Nervous System" was a desire to slip out of my own self-consciousness into direct, open experience of the world. In the right circumstances, the feedback loop of "Very Nervous System" effectively neutralizes consciousness, and can occasionally lead to states that could best be described as shamanistic. It can be intoxicating and addictive. I made a real breakthrough in the responsive quality of the system in 1987. I'd written a program where powerful drum sounds were produced by very aggressive movements. The result was extremely satisfying. After a week of developing and experiencing this new version, I found that I'd seriously damaged my back. I'd been throwing my body in the air with abandon, crashing myself against the virtual in search of those most satisfying sounds. This was a classic case of positive feedback.

Most natural and stable feedback systems are negative feedback systems, intended to keep a system in balance. If things get into any extreme state, the feedback mechanisms work against that state to restore balance and maintain the sustainability of the situation. This particular tuning of the "Very Nervous System" worked in reverse, egging me on to greater feats of physical movement until I wore myself out.

Pushing ourselves out of equilibrium is a way of opening us to change, but it can also lead to self-destruction or external manipulation. The mechanism that governs the evolution of life involves enormous test periods during which impossible or unsustainable life-forms are weeded out. Humans have evolved over a very long time to be well adapted to the stresses of everyday physical reality, and our species has evolved ways of balancing new pressures. But we now invent new pressures and stresses at an extraordinary rate. While technologies can be developed to counterbalance some of these stresses, the stability of this balance is not guaranteed. I'm not advocating a return to Darwinian rule, just pointing out the seriousness of the task of "engineering" this balance.

HAVENS FOR SAFE INTERACTION

The recent explosion of interest in interactivity surprises me. Interaction is so much a part of our daily life as to be virtually banal. Breathing is a profoundly intimate social and physical interaction: we breathe air into our lungs, extract oxygen, and expel carbon dioxide into the air, to be breathed in by others or transformed by plants back into oxygen. Talking, crossing the street, and driving a car are all interactions significantly more complex than those supported by most interactive computer systems.

The world with which we interact has becoming increasingly abrasive. We breathe in exhaust fumes. A growing list of foods interact with our bodies to cause cancer. Infectious diseases like AIDS make us squeamish about physical contact (whether justified or not). Under this bombardment, we're turning to ways of reducing our interaction with the world. The condom is, for example, a device intended to prevent interaction (either between sperm and ovum, or sperm and blood).

So perhaps the explosion of interest in interactivity is part of a search for havens of safe interaction: clean, sterile, non-physical spaces where we can satisfy our natural human desire to engage in things outside of ourselves.

THEORETICAL CLAUSTROPHOBIA

While the physical sterility of virtual experiences may be the easiest to grasp, the key type of sterility in artificial experience, for me, is that the ideas themselves are hermetically sealed. In "reality" our concepts, models and abstractions are always projections onto a complicated reality that never fully yields to our logic. Simulated experiences are built up from models that we have ourselves defined or already understand. In a contained interactive system, we enter into our own models, into a space of no true ambiguity or contradiction. There is no "unfathomable," which is a way of saying that there is no "God" in this virtual space.

SIMULATED COMPLEXITY

In a similar vein, it's important to understand the difference between "fractal" complexity and the complexity of life experience. Fractals are fascinating because a rich variety of forms are generated by a single, often simple algorithm.The endless and endlessly different structures of the Mandelbrot set are generated by a single equation addressed in an unusual way. This relationship between the infinite detail of the fractal and its terse mathematical representation is an extreme example of compression. The compression of images, sound and video into much smaller encoded representations is one of the keys of the current multimedia explosion.

Opposed to the incredibly compressible "complexity" of fractals is the complexity of true randomness. Something can be said to be random if it cannot be expressed by anything less than itself... that is to say, it's incompressible. This rather philosophical notion can be observed in our everyday on-line communication. To move data around quickly and efficiently, we compress it, then send it through a modem that compresses it further. What is left is the incompressible core of the information. As you can hear through your modem when you dial up your internet service provider, the result sounds close to random noise.

Randomness and noise are usually things we avoid, but in the purely logical space of the computer, randomness and noise have proven to be welcome and necessary to break the deadly predictability. But random number generators, used so often to add "human" spice to computer games and computer-generated graphics are not "random" at all. They merely repeats over a fairly long period?a sterile simulation of the real thing.

THE POWER OF RANDOMNESS

The classic story of the power of randomness is the story of the many monkeys at typewriters typing away for many years. The Laws of Probability suggest that one of these monkeys will at some point accidentally type the entire works of William Shakespeare. And if you accept evolution, then you could say that this has already occurred. Many sub-atomic particles working over a long but not infinite amount of time have managed to generate the works of William Shakespeare, by gathering quite arbitrarily into molecules, proteins, life-forms, social structures and ultimately into Shakespeare himself.

On the other hand, neither a fractal nor a pseudo-random number generator is capable of this feat. Those systems are "closed." No matter how far you expand them, Shakespeare's work will not be generated. Shakespeare is actually beside the point here. Replace the work of Shakespeare in the above discussion with any extremely unlikely but theoretically possible occurrence (the origin of life, the birth of the first consciousness or meeting the love of your life). These occurrences are statistically unlikely, but they can have a profound effect on the life of those who run into them. When you think back over your life, which were the really pivotal events: the predictable ones or the ones that seemed the most improbable?

In designing environments for experience, we must remain humble in the face of the power of irresolvable, non-fractal complexity. The computer is an almost pure vacuum, devoid of unpredictability. Computer bugs, while annoying, are never actually unpredictable unless this "vacuum" fails, as when the hardware itself overheats or is otherwise physically damaged. This vacuum is extremely useful, but it's no place to live.

When I started working with interactive systems I saw the "vacuum" of the computer as the biggest challenge. I developed "Very Nervous System" as an attempt to draw as much of the universe's complexity into the computer as possible. The result is not very useful in the classical sense, but it creates

the possibility of experiences which in themselves are useful and thought-provoking, particularly by making directly tangible that what is lost in over-simplification.

CONCLUSION

One of the initial motivations behind interactive interfaces was that they would allow users to apply their accumulated common sense and knowledge of the world to their navigation of the abstract realm of information. Abstract things become sensual and experiential. The use of familiar metaphors to approximate simulations of the real world enables users to make decisions and handle data in familiar, intuitive ways. This has been the reason for the dramatic success of the graphical user interface. In retrospect, however, this may have been merely a transitional strategy. Children now spend enough time interacting through synthetic interfaces that their common sense and knowledge of the "world" will have been formed partly by the interfaces and abstract simulations themselves. The shifting of the experiential base from "reality" to the video game's or educational software's virtual reality has far reaching implications. Interaction is not a novelty to today's children; it's an integral part of the only reality they have known.

From a purely practical point of view this is a useful situation. (Imagine a touch of sarcasm here) Children are adapting from birth to the language of synthetic interfaces. We will no longer have to worry about the real-world behaviours and expectations of our users that make designing intuitive interfaces so difficult. Common virtual sense will be widespread.

In the process however, those of us who design interactions inadvertently step into the realm of theologians and philosophers, perhaps even gods. We're laying the foundations for new ways of seeing and experiencing the world. And through communications interfaces we're building new social and political infrastructures. Economic pressure, intense competition, and shrinking product development cycles make it difficult to accept and do justice to this responsibility.

But accepting responsibility is at the heart of interactivity. Responsibility means, literally, the ability to respond. An interaction is only possible when two or more people or systems agree to be sensitive and responsive to each other. The process of designing an interaction should also itself be interactive. We design interfaces, pay close attention to the user's responses and make modifications as a result of our observations. But we need to expand the terms of thisinteractive feedback loop from simply measuring functionality and effectiveness, to include an awareness of the impressions an interaction leaves on the user and the ways these impressions change the user's experience of the world.

We're always looking for better input devices and better sensors to improve the interactive experience. But we also need to improve our own sensors, perceptions and conceptual models so we can be responsive to the broader implications of our work.

Notes:

1   quoted by O. B. Hardison, Jr., Disappearing Through the Skylight (New York: Viking Penguin), 319
2   Myron W. Krueger, Artificial Reality II (Reading: Addison-Wesley, 1991), xvi.
3   Florian Rötzer, "On Fascination, Reaction, Virtual Worlds and Others", Virtual Seminar on the BioApparatus (Banff: The Banff Centre, 1991), 102.