



*The
Design
of
Intelligent
Environments*

*Soft
Architecture*

*Warren
M.
Brodey*

WALKING TO WORK this morning I remembered the white-gloved policeman who is now replaced by computer-timed, radio-controlled traffic lights. All lights used to be of equal duration, regardless of the hour or the traffic load. They were "stupid" in the days before actual flow was fed back to change stop light duration. Flow projections and intelligent guessing are necessary features of our newer computer-controlled traffic systems: necessary for the speed and density of flow now common, for example, in subways.

Nevertheless, this intelligence of the subway system and a multitude of other similar computer-controlled systems is still like the automated control of a well run insect colony whose program for behavior leads them to compute approximately the same course of action repetitively, with little creative effort on their part to evolve a purposeful

behavior. When should this regulation which provides survival be called intelligence? I wonder if a man from the 17th Century looking at our present world would say that we had an intelligent environment. Would he be able to say that our environment was able to control itself more intelligently than his?

The concept of an intelligent environment softened by a gentle control which stands in place of steel bones and stone muscles is refreshing. A dam that senses impending flood and uses intelligence to prepare itself would not need be so ponderous. To date we have not endowed our environment with this creative flexibility; the intelligence we have commonly achieved is uncreative, stupid and in large measure hostile to human well-being. We have allowed hard shell machines to multiply and control us. Man is a captive of his increasingly automated mechanical environ-

ment. This process we have accepted ever since the early days of the industrial revolution, not imagining any other possibility. We have accepted the proposition that in order to use the power which machines deliver *economically*, we must restrict ourselves to the limited human behaviors that the machines can accept as meaningful control. One must steer by turning the steering wheel in the prescribed way regardless of one's body size, fatigue or personal style. Human behavior is mass produced by the power delivering tools man has learned to depend upon.

As we have created more and more power we have felt the iron gloves, which at first protected our hands from work, gradually thicken to protect us from touching the world around us. The teenagers search for a way back to "contact." But we cannot go forward by destroying the past. When man adapted for survival against a natural environment over which he had little control, he evolved; now men must evolve against the pollution of environment produced by our own progress.

What is the solution? Evolution now must include evolving environments which evolve man, so that he in turn can evolve more propitious environments in an ever quickening cycle. To stabilize the capacity we need to characterize this evolutionary dialogue. This characterization is increasingly being seen as the unsolved problem of our time. It is familiar to designers and architects in the student's question: "How do you design a house which will grow to meet the changes in the family that the house itself will produce?"

No man as yet knows the solution, but we can seek at least to clarify the question; a question well defined provides the beginning of its answer.

A DECADE AGO Rosenbleuth, Weiner and Bigelow wrote their historic paper, "Behavior, Purpose and Teleology." This ushered in cybernetic thinking. Their conception considered a thing and its environment in terms of their mutual relation. It defined behavior of the inanimate and animate within one frame of reference. The categories of behavior defined in that paper are a valuable start for developing a common notation for the design of intelligent environments.

Rosenbleuth, Weiner and Bigelow separated *active* behavior from passive behavior—behavior in which the object behaving is not a source of energy—as an object thrown. They subdivided active behavior into *purposeful* and *non-purposeful*. The latter is not directed to a goal, whereas the former is. If we decide, for example, to take a glass of water and carry it to our mouth we do not command certain muscles to contract to a certain degree and in a certain sequence; we merely trip the purpose and the reaction follows automatically. Although a gun may be used for a definite purpose, the attainment of a goal is not intrinsic to its performance. Some machines, on the other hand, are intrinsically purposeful. A torpedo with a target-seeking mechanism is an example.

In that historic paper the term "feedback" was first defined, and purposeful behavior was then separated into *feedback* or teleological and *nonfeedback* or nonteleological behavior. The word teleological was originally used to describe an innate or final divine purpose in all living things. The feedback control concept now allows us to define purpose without divinity: it is that goal from which deviation is corrected by feedback. The evolution of error correction procedures is used to define purpose; this brings us close to Darwin's concept of an evolutionary tree—a tree expanded in time by errors which escape correction and alter the feedbacks, but pruned by the death of those patterns which cannot survive when recontexted by the

evolving environment. Survival and purpose intermingle.

Feedback, or purposeful behavior, is in turn subdivided. It can be *predictive* or *non-predictive*. "The amoeba merely follows the source to which it reacts. There is no evidence that it extrapolates the path of a moving source. . . . A cat starting to pursue a running mouse does not run directly toward the region where the mouse is at a given time, but moves toward an extrapolated future position." Predictive behavior may be subdivided into different orders. "Throwing a stone at a moving target requires a certain order of prediction. The paths of the target and the stone should be foreseen. Prediction will be more effective and flexible if the behaving object can respond to changes in more than one . . . coordinate. The sensory receptors of an organ or the corresponding elements of a machine may limit the predictive behavior."

When the Rosenbleuth, Weiner, Bigelow paper was written, the existing automatic environments did not have the capacity to predict and extrapolate with sufficient complexity to be sensitive and responsive to self-organizing and evolutionary purposes. Given this capacity of our present machines, we can add to the list of behaviors defined in the paper. The category of the predictive machines can be further divided into *complex* and *simple*. An aggregation of simple machines grows only into a *complicated* machine decomposable into simple elements. The complex machine is more than an aggregate of its parts and their relations. It cannot be decomposed without destroying its capacity to maintain its organization. The complex machine can be further categorized as *self-organizing* (convergent) or *non-self-organizing*. In the latter kind of machine there may be sudden breakdowns, but in the former reliability is maintained by continuous breaking down and rebuilding. The system maintains its convergence by simplifying itself in terms of an internal purpose as defined by a complex net of intertwined feedbacks. If the self-organizing machine can maintain its purpose by responding to what was noise so as to evolve a new purpose, it can be called *evolutionary*. If it cannot, even though it is self-organizing, it is *non-evolutionary*.

GIVEN THIS HIERARCHY of behaviors of an object in relation to its environment, we can now redefine environment. Rosenbleuth *et al* defined it in these words: "Given any object relatively abstracted from its surrounding for study, the behavioristic approach consists in the examination of the output of the object and of the relations of this output to the input." When we speak of intelligent environments we traditionally define man as the object and the environment as the surrounding.

But we could also consider the surrounding as the object and man as the environment, or at least make them both object and environment to each other. Think of the effect of an infant's mattress or of his crib on the child. The balance of the mattress will affect the movement of the child—it will control him. It will teach him by subduing some movements and reinforcing others. The assemblage of rooms, walls and spaces in a home actively control the actions possible within it. An employee is trained by his work space and tools, a driver by his automobile. This concept of man as a passive unintelligent abstraction who does not create or evolve is a common simplification used by those concerned with environmental design. It is merely the reverse of considering the man as active and the environment as capable of only passive behavior. But much simplification is unwarranted. Imagine a time-lapse movie taken of a city and its inhabitants over the years. It would show an interaction involving purposeful, feedback, predictive,

self-organizing and evolutionary behavior. With finer grained measurements one could see evolution at work in a day or an hour. Learning itself is an evolutionary process—in its best form. It prunes out the obsolescent and allows the unknown to be realized.

Attending the recent Conference on Intelligence and Intelligent Systems sponsored by the Air Force Office of Scientific Research, I was impressed with their problem of teaching computers how to ascend the hierarchy of behavior outlined in the Rosenblyuth paper. These scientists know that their teaching must go through the evolutionary stages we see in the phylogeny of the creatures we know best—human beings. They provide the environment for their machines and teach them just as we are taught by our buildings and books. They build feedback and prediction into their machines and they are struggling to build in complexity, self-organizing reliability and evolutionary capability. Some say that a machine is not intelligent unless it solves problems without the help of its environment, without the help of a dialogue. But without being taught, man too would be unintelligent! As Marvin Minski puts it, in the *Scientific American* issue on Computers, we too easily measure the machine against a non-existent super-human man. Man must be continuously taught by his environment, both human and non-human. Man needs the novelty he metabolizes through his learning as much as he needs oxygen. To design intelligent environments we must know how to teach and are taught by our buildings, our work spaces, our transportation units. This process, being omnipresent, is easily unobserved. But we can caricature it back to our attention with our new control skills.

In the past the availability of energy was limited, and man's choices of what would be most pleasant to him as surroundings came after many compromises irrelevant to his creative survival and pleasure. Man has been a captive trained by the mediating devices with which he has controlled nature. But now intelligent environments capable of truly entering into dialogue are possible.

FOR THE SAKE of illustrating this trend and casting it into a form which reflects the new thinking, let me show a hierarchy of increasingly intelligent environments and the unsolved questions that prevent us from bridging the gap from complicated simple to truly complex, self-organizing and evolutionary designs.

An evolutionary environment maintains a hierarchy of long and short term purposes mediated by a complex network of feedbacks, each with its own dominant periodicities. These purposes themselves grow as the self-organizing system moves through levels of relative stabilization *vis à vis* its environment. The evolutionary environment is exemplified in the design of great art which grows in meaningful identity even though the perspective of the viewer is drastically changed by the impact of the changing information he draws from the art.

These respective levels of environmental intelligence are easily exemplified in practice, *until we seek to design-in the complex evolutionary dialogue*. As complexity increases, analytic logic undergoes what we might call graceful degradation—it slowly dies. We are then left with the relatively unformalized power of synthetic reasoning and simulation. We build crudely to find a way to think, and then we build again more exactly.

Some environments are *actively* intelligent, others *passively* so. A schoolroom that dampens the sounds of a creative surge of enthusiasm among its pupils in order to minimize noise transmission from room to room is *passively* intelligent. If it were actively intelligent, it would discrimi-

nate noise from nuances of speech, and it would have an active role in deciding what kinds of sounds were to be transmitted. It can be taught this intelligence by its designer. Active rooms may be purposeful or non-purposeful. The *purposeful* room, as a language laboratory—or a greenhouse—actively works to fulfill its purpose. It may have *feedback*—it may, for example, change its lighting and heat in terms of required plant growth. The sensors evaluate the plant growth as it takes place, and evaluate the resulting environmental changes. The motors modify heat and light in terms of this change. We are back now to the intelligence level of a computer-controlled subway, which senses passenger needs. Using the same model, let us consider the human teacher in the schoolroom who, like the traffic cop in his way, repeats the same window adjustment every time it rains or humidity increases. By the use of simple control devices, a factory or a schoolroom may be taught to change climate automatically. But such simple automatic devices have only the crudest feedback system. They do not actively sense worker efficiency, as a function of climatic change.

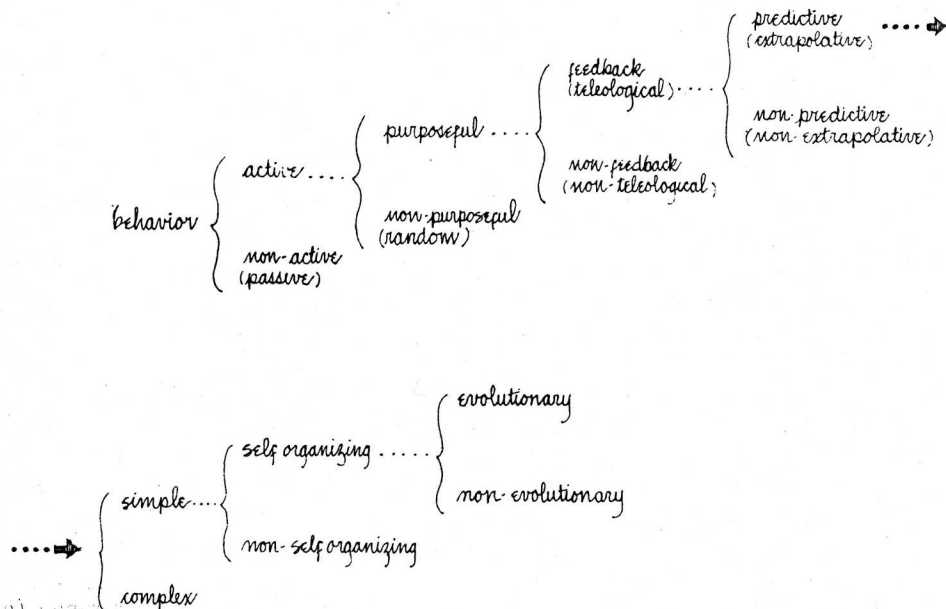
PRESENT CONTROLS are like those of an adding machine that pays attention to the user only through the commands it is given. It regulates the human being by making all but his simplest, most ritualistic commands meaningless. Environments that use more feedback related to user-machine designed purpose can be *predictive* of likely change, given the changes that have already taken place. It can be predictive in fewer or more dimensions depending on its elegance of communication with its context (one or more sensors).

An environment may be simple or *complex*. If simple, it can still be made complicated, but multiplication of simple people or of simple devices does not create intelligence, unless they are organized onto the next level—a complex group.

I recently visited an expensively automated home where the gadgets were beginning to turn each other on and off by mistake. The ring of the telephone had the same frequencies as the T.V. channel changer—and changed the channels. In a complicated “deluxe” automobile of a friend, the gadgets all interfere with one another and will occasionally go on a rampage, wildly clicking each other on and off. The doors are designed to lock automatically at 8 miles per hour, the lights to dim at an approaching light, the heater to air-condition, and on on. Now the approach of lights locks the doors—the driver swears that the gadgets are trying to build feedbacks between them.

If the simple gadgets were actually to build a network of self-reinforcing feedbacks between them they could then be controlled, not only by changing individual readings but by shaping their beginning *self-organization*. What do I mean by self-organization? How does one help a number of individuals form a group? By discovering a purpose where three elements at least cannot be simply divided into two without loss of function. We wish to use, for characterizing the process, those aspects of complexity which evolve as the simple system made too complicated goes wild beyond simple control and starts to organize itself.

Why try to design a complex environment? Present control sophistication allows us to construct a learning environment in which lighting, heating and information display can all be usefully controlled in an intelligently interconnected way and corrected by pupil preference expressed in voted nods of “please open the window—let in more air.” The teacher who respects individual variation and children's creativity can develop a more evolutionary style.



But she often falls victim to fatigue. The non-human system does not allow her to simplify control or to predict. There may be no windows; the thermostat is set for the ideal child. All children must accommodate to averaged environments built primarily to reduce upkeep and to last for many generations. Children are soon taught by the school environment that learning means disregarding personal variations; the captured child must adjust. He is forced to adjust *not* by the school or the teacher's making a decision, but because his individual variations are meaningless within the range of values allowed by the environment. Children cannot be creative in an environment of paper forms, bricks, mortar and air-conditioned hums whose unintentional purpose is to reduce all to an average which stifles the drive for discovery and change. Those who have tried to deal intelligently with true-false questionnaires constructed for simplified data processing know the "stupidification" that results. The teachers and the school system, as well as the child, are made stupid by the complex media and controls at present used. An environment that did not need to simplify children into square feet of space per child would allow them more aliveness. An environment that would learn from each child his style and help him to evolve it would be a true learning environment. It would not be just a caretaker—it would take part in his evolution.

MORE INTELLIGENT environments are now being developed—the Apollo space ship system is one example. But though the space ship environment is a masterpiece of technical achievement, it is still only a *complicated* environment: it is neither complex, nor self-organizing, nor evolutionary. It survives the loss of a component only through a reserve of back-up components that must be themselves reliable and that must be switched into place by equally reliable redundant links. As such essentially stupid environments become more complicated, dials and toggles soon stand in massive array. All the skill of human engineering is required to avoid the mistakenly flipped switch that at supersonic speeds spells sure disaster.

The bottleneck is now the lack of an intelligence match between man and computer, for man cannot use his evolutionary skill when the environment has none. The ideal environment would replace toggles and switches by a skill-

ful mutual man-machine sensing of the advantages and disadvantages of a particular cooperative behavior. The environment system would itself grow with the user.

ARE SUCH COMPUTER BASED intelligent environments possible? We have developed our environment through the stages of being active, to having purpose, to using feedback, to extrapolating, to being complicated. Can we now teach our machines or environments first *complex*; then *self-organizing* intelligence which we can ultimately refine into being *evolutionary*? The first answer to this question is that we do not know how.

But if we do not know how to create a complex self-organizing evolutionary environment, we can at least begin assembling the ingredients and concepts that can, by trial and error, produce better tools. If the task is impossible we shall at least learn why.

A *complex* system will include a complex network of interconnected feedback loops. What example could we use to start with? I would choose to place a man and his computer-controlled environment in close connection, so that the man who is *in* the evolutionary exploration process can tell us how to sharpen our exploration. Let us put the man in a room in which the air flow, temperature, lighting intensity and color, the acoustical reflection and conductivity, the floor vibration and as many other known parameters as possible will all be computer controlled and measured. This has been done, to some degree, but not all in one place or at one time. The man will also be instrumented so that his behavior can be monitored. We will use as many ways as we can of measuring the man's outputs—both physiological and behavioral: heart rate, electroencephalogram, surface heat, core heat, head movement, hand movement, etc. These data were used in the past to describe what man is like; we will be satisfied if we can help one man to evolve a meaningful learning dialogue with his personally designed environment. Now, remembering that our purpose is to develop a truly interconnected network of feedbacks, let us connect the man's output behaviors—heart pulse accelerations, for example—so that they become data which the computer uses to adjust environmental parameters. Let us make these connections so that each is connected to all the others. Now we have a complex network of feedbacks—we can no longer tell, in traditional terms, which is man

and which is machine. But our purpose is to simulate a complex system—or at least to build a caricature of it that will help us to learn whether such a system can be built.

Our next task is to see if this complex system can become *self-organizing*. We have put a man in the loop. The information of the system flows through him as a part of the organizing network. If his heart beat accelerates, the room becomes redder (for example); if his breathing deepens, the room takes on a richer hue. As the hue intensifies his heart may beat faster in response to the stimulus (the strength of color which changes with his feeling). This personalized total environment will be capable of producing a profound experience without brain damage. If the eyes move to the left the display may adjust to the right, or become dimmer because his heart has “reddened” the room. The computer will be taught to use these extrapolations of its data most suitable for providing an experience and pattern that man and machine can organize.

Let us ask the man if he can discover any patterns in the system, patterns which he can try to organize. As he learns he will expand what he has already begun to know. Perhaps in the accelerating dialogue he senses a sudden rhythmic beat reminiscent of a jazz band. Let us help him recreate and measure that beat. We will shape the system by changing its sensitivity (the amount of change necessary to get a message through the net) or the time delay (the time it takes to get a message through) or other overall system features. We can link the machine inputs and controls to the computer’s memory of the system’s state when the jazzlike occurrence happened. Can the man recreate the old stability? If he can learn to use his many outputs, now computer inputs, to stabilize the complex environment—if he can change the computer program so that it learns to join him in this effort to find an easily recognizable pattern—then we will have begun to study the process of self-organization. We will have begun to understand that convergence of data necessary for maintaining a complex organism so that it changes noise into information which allows the system to stabilize even as it changes. Human beings organize their human environments this way with ease—a child and mother must do it or a household will never settle down. We are trying only to simulate a common occurrence.

Having created such an artificial man-machine system—a soft environment—let us now confront the man with the task of *evolving* a time phrased purpose with many kinds of goals, interrelated in time with the others. This is not essentially different from the problems facing a woman who prepares dinner while looking after children, paying bills and answering the phone—with only a small amount of the information processed ever being thought through consciously. If asked she cannot analyze her intelligent procedures. She corrects for error, she pays attention only to

what needs attention. She knows when the family has organized itself and found a purpose for the day. She groups data and tasks, constantly changing the code to help the job organize itself. If you have to ask a question unrelated to her state there can be no answer. She is aware of her evolving behavior only when her rhythm is broken.

ALL HUMAN BEINGS depend on their self-organizing and evolutionary relationship with their environment. A man automatically changes his voice when he enters a new room so that it will still sound like him. But can there be an evolutionary stage of the postulated system? The evolutionary stage of the self-organizing complex man-environment system can only grow out of its antecedents. Even human beings as evolutionary creatures can only develop each new refinement of being out of the last. We cannot do better.

WHAT WILL WE ACHIEVE if we can build the kind of intelligent, complex, self-organizing and evolutionary system that I propose? We have new tools to try out. We have real time computation available. A computer can predict from the trajectory of a man’s hand where it can go next, and within this range (given the man’s purpose as demonstrated by his last movements) where it is likely to go next. It can use this prediction to make available to the hand the implements or the light it may need. The computer can discover that a child has not been paying attention because he is bored, once it is taught the particular behavior patterns that indicate boredom. The teacher and the student would both teach it intentionally and perhaps unintentionally once they decided that this was their wish. The computer could change the fresh air in the room, or the lighting or the lesson, or the size of letters, the mix of spoken words, pictures and alphanumeric, or color, or two or three dimensional display. It could request that the teacher appear when the child-computer system encounters difficulties. Another analogy would be the dynamic transit system which maintains its purpose in relation to the town it serves even as it and the town change through their efforts to maintain equilibrium.

This design of intelligent environments, this idea I call *soft architecture*, as yet seems to interest few professionals. But the increasing capacity and lessening cost of new computers offer us the tools now. With the flux produced by our explosive progress, if we can begin only by modifying the school environment so that it actively teaches children, we will at least see the next generation taking highly intelligent environments for granted. Our progress will depend on those few who are willing to accept and apply these new concepts. Limitless opportunities for applying the new control sophistication will appear, once we recognize as obsolete the old economic pressures which reduce people to that average required by a rigid external environment—once hard architecture begins to be replaced by soft.

