

1-1-2001

## Intelligence and Autonomy

Ted Krueger

Follow this and additional works at: <https://newprairiepress.org/oz>



This work is licensed under a [Creative Commons Attribution-Noncommercial-No Derivative Works 4.0 License](https://creativecommons.org/licenses/by-nc-nd/4.0/).

---

### Recommended Citation

Krueger, Ted (2001) "Intelligence and Autonomy," *Oz*: Vol. 23. <https://doi.org/10.4148/2378-5853.1361>

This Article is brought to you for free and open access by New Prairie Press. It has been accepted for inclusion in *Oz* by an authorized administrator of New Prairie Press. For more information, please contact [cads@k-state.edu](mailto:cads@k-state.edu).

## Intelligence and Autonomy

*Ted Krueger*

The terms “Intelligent Buildings” and “Smart Buildings” have been used in the building industry for almost two decades. They refer to a wide variety of capabilities integrated into the building fabric based on computer and communications technologies. One of the truisms of the late twentieth century is that the development of microprocessors has had a profound effect on the culture. The availability of commodity desktop and portable computing machines and the networks that connect them has altered the ways that we work and communicate. Yet despite their ubiquity, and their role in both funding and driving the development of chip production, the number of microcomputers in use is far exceeded by micro-controllers that are embedded into objects in the environment.

Computer technologies not only allow for the transformations in the way we think and communicate but have begun to alter the fabric of the environment that we inhabit. We may find that the most profound changes due to the introduction of computer and communications technology in architecture come not from the availability of visualization, presentation, and documentation tools that they provide, but from the incorporation of intelligence into the materials and systems from which we build.

As Hunt<sup>2</sup> notes, the cybernetician Gordon Pask understood the rela-

tionship between the human and an adaptive and interactive architecture more than three decades ago. Pask posited that the domain of design was not the determination of the form of the building but the structuring of the social context in which humans interacted with their environment and with each other. He used the term “mutualism” to designate a kind of symbiotic relationship between the architecture and its inhabitants.<sup>3</sup>

Architects have in the interim focused on the symbolic and linguistic dimensions of their work and on formal gymnastics rather than architectural behaviors leaving the work of integrating intelligence into architectural media largely to manufacturers, property developers, and facility managers. These entrepreneurs have defined the problems to be solved and have specified the solutions to be implemented. They have also determined what is commonly understood as “intelligent buildings” with all the elasticity that marketing engenders. This paper will review some of what is considered to be “intelligent” in architecture and will review contemporary theories regarding the nature of artificial intelligences, in an effort to find common concepts and implementations and to begin to discern the general direction in which a more disciplined notion of intelligence may lead.

The aim of much contemporary intelligent building research is to find

technological solutions that promote the competitiveness of the construction industry and its clients.<sup>4</sup> In these contexts, an “intelligent” building is one that affords “productivity and cost effectiveness by means of optimally designed and interrelated structures, systems and subsystems, and services and management.”<sup>5</sup> The purpose of these integrated and centralized systems is an economy of building construction, operation, maintenance, and administration. The objective is efficiency, optimization, and control.

### Intelligent Subsystems

Uses for embedded intelligence in architectural applications began in the late 1970s with environmental control systems<sup>6</sup> and occurred at the intersection of a reduced cost for the control technology and rising energy prices. The use of micro-controllers for the management of building subsystems—such as security, fire protection, environmental controls, vertical transportation, and lighting—is driven by the increased capability that they made possible relative to the available electrical and mechanical controls. This led to greater operational efficiency, flexibility, improved occupant comfort, and reduced energy and maintenance costs. Throughout the following decade, manufacturers developed a wide variety of proprietary control systems, each optimized in isolation from the other systems installed within the same structure.



Embedded computing derives much of its functional advantage from the ability to integrate multiple inputs into a single behavior—that is to produce behaviors that are specific to a range of conditions. Computerized control techniques allow for complex relationships to be accommodated not only within a single subsystem but also, through intercommunication, with a variety of related subsystems as well.

The fire protection system in a building employs smoke and thermal sensors to detect an emergency and activate the sprinklers or other fire suppression devices. It also must alert the fire department, sound local alarms, return elevators to the ground floor for use by the authorities, pressurize the escape stairs to eliminate smoke infiltration, turn on exhaust fans to evacuate smoke from the building, and provide for emergency exit lighting. This example, though extreme, illustrates the benefits of integrating a variety of building subsystems. The building industry is concerned with

the provision of many disparate but interrelated systems that must be integrated into a structure. Methods of coordinating between them promises to result in benefits equivalent to those that drove the initial incorporation of computing devices.

Intercommunication methods were devised, integrating a full range of intelligent subsystems into unified whole. The first step in rationalizing this integration was the development of a single communication cable—"structured cabling"—unifying security fire alarm data and telephone networks with the traditional building management systems.<sup>7</sup> Integration implies common standards for communication hardware and software and protocols for the communication between devices. At present, these standards are newly adopted<sup>8</sup> or under development.

The provision of communication and data services over local and wide area networks is generally considered a marketing feature of the intelligent building. Due to the rapid technical development that is occurring in this area, these services must also be frequently altered, upgraded, or replaced. Providing for these periodic modifications may be considered intelligent design, but does not increase the inherent intelligence of the building in which they are incorporated. The incorporation of these services is essentially an independent system.

With structured cabling, however, these networks serve the intercommunication requirements for building services and subsystems. This common communication infrastructure enables contact between buildings within a complex as well as with other structures within the district<sup>9</sup> or, in the case of Singapore, communication across the whole of a city-state.<sup>10</sup>

It is clear that the implementation of intelligence in the architectural context has occurred in response to a desire to increase the functional capabilities of a variety of subsystems. There was, initially, no overall development strategy and little coordination of effort. Presently, the development of standards demonstrates an interest in intercommunication, interoperation, and an effort to open development to a wider range of commercial sources.

### Intelligent Materials

In addition to developing automated versions of systems that are commonly used in architectural settings, there is an interest in devising strategies for increasing the intelligence and capabilities of the materials from which we build. Substantial work in intelligent materials has been undertaken with applications intended for the aerospace or defense industries. The strategy employed and the methods used have evolved directly from composite materials—the typical method of fabrication used for intelligent structures. A range of materials,

each with a desirable property, are brought into the proper orientations and fused in a thermoplastic matrix. For example, a variety of fibers—glass, graphite, and ceramic—may be oriented to directly counteract the forces that are anticipated on the component, their properties individually matched to the anticipated stresses. Optic fibers are placed as sensors to track the forces impinging on the structure as well as its internal states during manufacture, installation, and use. These sensors provide information to embedded micro-controllers that record and process it. Composite materials may include actuators, such as shape memory alloy wires, to counteract forces or damp vibrations as directed by the embedded controllers.<sup>11</sup> Composite techniques result in

a material with enhanced structural properties, material that has the ability to sense, record, decide and react.<sup>12</sup>

Efforts to integrate sensory components into architectural and civil structures are in their infancy. Reinforced concrete, the architectural equivalent of composite materials, has received considerable attention.<sup>13</sup> Houston and Fuhr at the University of Vermont have embedded fiber optic sensors on a highway bridge, a hydroelectric dam, and a university bio-medical building and are collecting the resulting data.<sup>14</sup> This information will inform analytic models and promises to provide an empirical basis for structural theories. The availability of sensory technologies may make it unethical to build structures that do not warn users of



impaired capacity or impending failure.<sup>15</sup> Sensory and reactive techniques applied to the substance of a building represent a fundamental increase in capacity beyond that offered by the control of integrated sub-assemblies. Safety, economy, and efficiency will drive the incorporation of intelligent materials.

### Behavior-based AI

Within the context of the present paper, it is reasonable to turn attention from the details of the integration of intelligence in buildings to a consideration of the methods that are under development in the field of artificial intelligence. These may provide information about the general direction that this integration may take.

The need to undertake activities that were by their nature difficult or tedious for humans to perform—cryptographic analysis, calculating projectile trajectories and sorting census data, for example—drove the early development of computational devices. Initial efforts in the field of artificial intelligence focused on related problems of logic and symbolic processing. In these areas, artificial systems made rapid progress towards the goal of equaling and surpassing human capabilities. Robots, as a platform for investigating intelligence, develop as an interface between a symbol processing system and the environment. The robots iterated through a sense, model, plan, and act (SMPA) paradigm developed out of the human experience of intelligence that was validated through introspection. Decisions were made by reference to rule structures developed by the researchers and embedded in the controlling programs for the robot.

The machines were slow and clumsy in their interactions with the world. Extensive research and computational resources were expended on the problems of developing a comprehensive world model from sensor input and in planning appropriate behaviors

within that world. Paradoxically, the more information that was gathered by the robot the more hesitant were its interactions. A “paralysis through analysis” resulted from the explosion of possibilities intrinsic to the modeling and planning activities.

Inherent in SMPA is a Cartesian mind/body dichotomy implemented on an electro-mechanical apparatus. The emphasis was clearly on the logical and symbolic operations—the mind—with the robotic—the body—as a secondary and subsidiary operation within the hierarchy. Searle’s critique of artificial intelligence indicated that the symbolic approach lacked a means by which the symbolic tokens could correlate with the kinds of experience from which they could derive meaning.<sup>16</sup>

An alternative approach was based on a formulation of intelligence that was measured by the capacity for appropriate and timely behavior of the robot within the “real world.” Simple independent task-achieving modules were written that allowed the robot to interact in its context. Once tested and debugged, these modules remained in place and attention was given to writing higher-level structures that would be concerned with the mediation and coordination of the simpler behaviors in an effort to produce more complex interactions with the environment. Viable control systems result from each layer. A module that directed the robot to move forward, for example, might be subsumed by a higher-level program concerned with input from the sensors determining the location of an obstacle to be avoided. That program might itself be subsumed by one which planned a route from one location to another and so on to arbitrary levels of complexity.

This hierarchically layered structure, developed by Brooks and others, was termed Subsumption Architecture. Distinguishing characteristics of this approach are its modularity, robustness,

embodiment, situatedness, and capacity for emergent behaviors.<sup>17</sup> Because each behavior-generating module operates independently, mechanical or behavioral failures affecting higher-level behaviors leave the robot with its more primitive capabilities intact. The behavior of the system is robust, rather than brittle, in the face of unforeseen circumstances. Rather than focus on the abstract processing capacity of the robot, at the most basic levels, these machines were built on sensor-effector couplings; they are inherently embodied and situated within a context. The control of the machine is not centralized but may be understood as a system of independently acting agents.<sup>18</sup> In this way, it is related to Minsky’s “Society of Mind,”<sup>19</sup> an agent-based understanding of human intelligence.



The strong relationship between the machine and its environment allowed for the development of emergent behaviors. To take a simple example, if a module directing the robot to move to the left were subsumed by a module that directed it to move to the right for a while after encountering an object, the robot would follow a wall—by periodically bouncing off of it—down a corridor and, upon encountering a door, move through it and continue to follow the wall on the opposite side. In this case, the wall following


behavior is not programmed into the machine but emerges by means of an interaction between the machine and its environment. This aspect of the robot’s behavior is in many respects similar to the flocking activity generated by Reynolds. In that case, the activity of the group was an emergent property developing out of simple rule structures governing the behavior of individual agents and their context.

This behavior-based approach to artificial intelligence was designed to follow a development reminiscent of biological evolution. Rather than starting with human level competence as an objective, the capacities of simpler organisms were initially modeled, and more complex behaviors were built gradually over time.

### Architecture and Subsumption

The transfer of notions of intelligence from robotics to architecture may initially seem inappropriate. Architecture is static in comparison to a mobile robot. One of the foundations of architectural education is the study of statics—an analytic tool of structural analysis. Society expects and requires a static architecture. A building with any significant dynamic is a prime candidate for a lawsuit. If one examines the situation from the standpoint of the sensory apparatus, however, it will be realized that the sensor responds only to a change of state. It has no *a priori* notions regarding movement. Therefore developments founded on mobility will be applicable in any context where there is a significant dynamic. In architectural contexts, this dynamic may be the result of programmatic as well as environmental or structural variables.

The subsumption approach is applicable to architecture in several ways. It is an incremental and iterative strategy that is capable of yielding useful behaviors at frequent points in its development. Complex behaviors are built over time on the foundations provided by simpler



pre-existing ones. There is no need to design a complete system from the ground up—higher-level behaviors assume the existence but not the authorship of simpler ones—nor to integrate all systems into a unified network, but only to establish an asynchronous intercommunication between them as the need arises. Some of the lowest level sensor-actuator couplings are already developed and in place in mechanical, lighting, and security systems, and other more sophisticated systems exist within most appliances. Standards have been devised that allow these devices

to intercommunicate over networks. Architecture is inherently embodied and situated within a dynamic environment. This sets a context where higher level behaviors can be developed. The useful behaviors resulting from each layer of integration provide a motive for an increase in sophistication. An ideology of utility and the entrepreneurial imperative in contemporary economics guarantees this kind of development. A further advantage is that the system degrades gracefully if there is a failure - the lights stay on and the elevators still work.

Viewed from a chronological perspective, the implementation of embedded intelligence in architectural settings has paralleled the subsumption strategy used on robotic platforms. From the initial use of micro-controllers in isolated subsystems, to the development of intercommunication protocols that allow for the integration of functions across subsystems, these developments have not occurred by reference to an overall strategy but as a pragmatic response to functional issues and an effort to generate useful behaviors. A further consideration of the disciplined study of intelligence may illuminate the potential limitations of this line of development, because behavior-based approaches are not without difficulty. It is unclear how well the reactive strategy will scale up in complexity. Maes notes that emergent behavior, even at relatively simple levels, requires that the designer set an interaction loop between an agent and the environment that results in a convergence on the desired behavior.<sup>20</sup> In this case, emergent behavior is not fortuitous, but a function of the effort and skill of the designer. With increasing system complexity, the number of behaviors and the probability of conflicts between them increases. The programmer must anticipate a combinatorial explosion in the set of possible interactions and develop the means to address them.

Architecture as an extended object with a large number of behaviors running concurrently is a prime candidate for this kind of complexity. One of the fundamental concerns of design—architectural, robotic, or otherwise—consists in the mediation of conflicting requirements. These conflicts carry through beyond design into the operation of the artifact. Technology may supply new means to address—but not eliminate—these problems.

Harnad suggests that, to produce an artificial mind that could be equated with a human one, a method of ground-

ing its symbolic capacities must be employed.<sup>21</sup> This possibility arises when the symbols are organized relative to contact with the world and their interrelations are established through experience rather than by pre-processed rule structures. The artificial mind must be embodied if it is to escape Searle's critique.

Brooks acknowledges that the behavioral and symbolic approaches may be complementary.<sup>22</sup> While the extent to which behavior-based systems can integrate and ground symbolic ones remains an empirical issue, the deepest levels of integration will arise from situations in which the representational systems are derived from experience. They may have a development and structure that is closely related to the behavioral parameters from which they originate. It is in this sense that contemporary theories of artificial intelligence address Searle's critique. The emphasis on embodiment and situatedness provide a context on which the derived symbol tokens can be grounded. The relationship between the symbol and experience is what allows the symbolic to have a meaning.

This paper has not considered all uses of media in architectural settings but has focused on basic functionally-driven implementations of computer and communications technology. There are projects that have used media in their visual and expressive capacities as elements of a formal system. "Mediarchitectures" as theorized by Thompson represent a potentially important alteration of the role of the architectural object within a cultural dialogue. They are not considered here because there is as yet no link by which the text and images that are used as elements of a facade, for example, could have import to the cognitive dimensions of the building. This relative isolation is also true of the data which flows across the telecommunication networks embedded within a building. The presence of data or symbols does

not guarantee its saliency even if the structure participates in its processing or transformation. What can be understood from this research is that these signals must have a relationship to the embodiment of the system if they are to be relevant to the intelligence of the system. In this case, the rate of data flow in a network may be important to the structure; the specific content of that flow may not.

### Robotic Imagination and Desire

Stein has undertaken some development that is particularly intriguing relative to embodiment and intelligence. In her study of memory and imagination in a behavior-based robot, the machine was capable of understanding places it had encountered by virtue of a signature 'feel' or representation of the place as encoded by its sensory apparatus. These perceptions were stored for future reference. The researcher was able to feed the robot a representation of a place it had not visited but which it could compare with its current sensory constellation. It was found that the robot could explore the world and recognize places that it had not previously visited but which it had been told about and for which it had, in a sense, an image in its imagination.


While it has been suggested that the logical and symbolic capacity of the human mind is a serial process riding on top of the parallel structure of the unconscious, Stein's approach was not to place the representational capacity atop the behavior-based substrate, but to run the two in parallel. The physical structures of the machine were emulated on one of the robot's computers and the robot could explore these virtual spaces using its simulated sensors.

The existence of the parallel structure enables a capacity for speculation. The machine's current sensor readings could be transformed based upon hypothetical grounds and evaluated with

its virtual sensors. Several scenarios could be produced and evaluated for the best course of action.

At this point, the question of the basis of those evaluations will arise simultaneously with the issues of which scenarios to generate and how many to produce, as well as how to evaluate of the number of steps into the future to continue the exercise. The constraints faced by the responding system must be anticipated either by a programmer or by rule structures imbedded in the agent. Here again, we are again up against the combinatorial wall. The range of permutations may be circumscribed by reference to goals.

Goals not only bias or restrict the range of choices but also make possible the allocation of resources such that several simultaneous objectives may be addressed or approached through a single action. Typically, goals are not isolated, but rather occur as a nested series of interrelationships that provide a guide for action at many scales simultaneously.<sup>23</sup> As Maes notes, goals are a crucial ingredient of self-consciousness and a precondition for effective learning. One of her contributions to research in this area has been to develop autonomous agents that are capable of setting their own goals in the face of conflicting requirements.



### Intelligence and Autonomy

There is a common developmental path that can be discerned within this research. Increases in complexity cannot be met with brute-force

programming. The requirement that all possible states of interaction be anticipated and provided for sets a practical limit on what can be achieved. These limitations are simply a matter of mathematics. Repeatedly, as the combinatorial limitations become manifested, the designer is required to develop some means by which the agent can independently evaluate its context and take action. A measure of autonomy must be granted to the machine in order to be able to deal effectively with the complexity of its interactions with the environment. Autonomy is a fundamental change in the nature of the artifact that in turn requires a re-evaluation of roles that objects play in both the cultural and cognitive processes. It is this aspect of embedded intelligence that most profoundly alters our relationship to the products of our material culture.

Agre conveys the change that the principled characterization of interactions between agents and the environment has had for the field of artificial intelligence.<sup>24</sup> The concept of agent may embrace human and other biological entities, machines, and software. Each may be to some degree autonomous, intelligent, and situated within an environment, their roles interlocked and complementary.

Huberman and Clearwater have implemented a multi-agent system for controlling the thermal resources of a building at Xerox PARC.<sup>25</sup> In this scheme, computational agents linked directly to temperature sensors and air-flow actuators bid to buy or sell resources in an auction. This market-based system has provided a more equitable distribution of resources than has been available using conventional control schemes. The agents, in this case, act with completely autonomy beyond the thermostat settings established in each office by its human occupant. Agents trade in a virtual money established for the auction. As this capital has no direct corollary in

the world of the occupant, the authors note that it is possible to skew the results of the auction by setting the thermostat unusually high or low. If the humans were not isolated from the dynamics of the control system, that is if the virtual costs became real, the resulting interlock of roles would return the system to stability.

This interdependence suggests that the mutualism suggested by Pask may in fact be grounded in the findings of contemporary research on intelligence. Architecture in this context becomes symbiotic with its inhabitants.

As a discipline, architecture is oblivious to the possibility that autonomous intelligence may develop and so may be unable to recognize it if it occurs. As there may be no direct formal implications, many architects will not be interested and need not concern themselves. However, given some probability that intellect will arise, there may be those who would attempt to engage it and to take a hand in its development. There is a strong relationship between the structure of developments in intelligent buildings and the strategies employed in artificial intelligence. It is tempting to expect that some non-trivial forms of intelligence may emerge from these processes. Yet, as noted above, emergence is no magic. If intelligent environments are to occur, they will do so by virtue of the craft of their creators.

An interactive and adaptive architecture indicates that the locus of design migrates from form to the parameters of behavior. Its intent shifts from control to facilitation, from a restriction to an amplification of the design space and a consideration of the interface that will be required for intelligent and interactive environments.

## Notes

1. "Intelligence and Autonomy" was first published in *Convergence, the Research Journal in New Media Technologies*, 5, no. 1: 91–101. This article is an elaboration of the paper "Autonomous Architecture" presented to and published in the proceedings of *Consciousness Reframed: Art and Consciousness in the Post-Biological Era* by the Center for Advanced Inquiry into the Interactive Arts (Newport, Wales: University College Newport, 1997) and in a special issue on "Computers and Post-Biological Art" in *Digital Creativity*, 9, no. 1.
2. Gillian Hunt, "Intelligent Architecture," in *Consciousness Reframed: Art and Consciousness in the Post-Biological Era: Proceedings of the First Conference of the Center for Advanced Inquiry into the Interactive Arts* (Newport, Wales: University College Newport, 1997).
3. Gordon Pask, "The Architectural Relevance of Cybernetics," *Architectural Design* (September 1969).
4. Mervi Lehto, "Intelligent Building Research," 1995. Manuscript available at: <http://www.vtt.fi/rte/ais/projects/rtemle/alytaeng.htm>.
5. Gordon Friedlander, "Smart Structures," *Mechanical Engineering* (October 1988): 78–81.
6. Marc Harriman, Marc, "Buildings Get Smarter," *Architecture* 80 (June 1991): 107–110.
7. Paul Nicholson, "Towards a Digital Janitor," *The Architects Journal* 201 (June 1995): 38.
8. Steven Bushby, "BACnet: a standard communication infrastructure for intelligent buildings." *Automation in Construction* 6, Elsevier Science B.V, 529–540.
9. H. Allison Smith, Yoshitaka Takeuchi, and Hareesh Shah, "An Advance Notification System for Smart Structures in Seismic Zones." in *The Proceedings of the 1st European Conference on Smart Structures and Materials, Session 7*, Glasgow.
10. Juzar Motiwalla, Michael Yap, and Lek Heng Ngoh, "Building the Intelligent Island," *IEEE Communications Magazine* (October 1993): 28–34.
11. Eric Udd, "Fiber Optic Smart Structures," in *Proceedings of the IEEE* 84 (January 1996): 60–67.
12. Edward Crawley, "Intelligent Structures for Aerospace: Technology Overview and Assessment," *ALAA Journal* 32 (August 1994) 1689–1699.
13. A. Mendez, T.F. Morse, T.F., and F. Mendez, "Applications of Embedded Optical Fiber Sensors in Reinforced Concrete Buildings and Structures." *Fiber Optic Smart Structures and Skins II, SPIE*, Vol.1170 (1989).
14. Dryver Huston and Peter Fuhr, "Intelligent Materials for Intelligent Structures," *IEEE Communications Magazine* (October 1993) 40–45.
15. Raymond M. Measures, "Smart Structures with Nerves of Glass," *Progress in Aerospace Science* 26 (1989) 289–351.
16. J. Searle, "Minds, Brains, and Programs," *Behavioral and Brain Sciences* 3 (No. 3, 1980): 417–457.
17. R. Brooks, "A Robust Layered Control System for a Mobile Robot," in *A.I. Memo 864* (Cambridge: M.I.T. Artificial Intelligence Laboratory, 1985), 1-25.
18. Rodney Brooks and Jonathan Connell, "Asynchronous Distributed Control System for a Mobile Robot" in *Mobile Robots*, SPIE 727, 1986.
19. Marvin Minsky, *Society of Mind* (New York: Simon and Schuster, 1985).
20. Pattie Maes, "Designing Autonomous Agents," in Maes, ed., *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back* (Cambridge: Bradford/MIT Press, 1991) 1.
21. S. Harnad, "Does Mind Piggyback on Robotic and Symbolic Capacity?," in Morowitz and Singer, eds., *The Mind, The Brain and Complex Adaptive Systems; Santa Fe Institute Studies in the Science of Complexity, Volume 22* (Reading, Massachusetts: Addison Wesley, 1995) 203–220.
22. R. Brooks, "Elephants Don't Play Chess," in Maes, ed., *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back* (Cambridge: Bradford/MIT Press, 1991) 3–15.
23. Pattie Maes, "Situated Agents Can Have Goals," in Maes, ed., *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back* (Cambridge: Bradford/MIT Press, 1991) 49–70
24. P. Agre, "The Soul Gained and Lost, Artificial Intelligence as a Philosophical Project." *Stanford Electronic Humanities Review. Constructions of the Mind* Volume 4, Issue 2 (June 1995).
25. B. Huberman and S. Clearwater, "A Multi-Agent System for Controlling Building Environments," Xerox PARC manuscript of a paper presented at ICMAS-1995 available from the authors.