Information Architecture and the Emergent Properties of Cyberspace

Andy Norman
Peter Lucas
MAYA Design Group, Inc.
Norman@maya.com

> "The proliferation of microprocessors and the growth of distributed communications networks hold mysteries as deep as the origins of life [and] the source of our own intelligence…" –George Dyson

Some say the science of complex systems had to await the advent of computers. If this is true, these sciences owe an intellectual debt to the designers of sophisticated information technologies. Can systems science repay this debt by helping engineers understand the system-level implications of basic architectural design decisions?

Information Systems

The most striking characteristic of our age is the rampant proliferation of devices that process information. In 1997, the number of microprocessors in the world reached six billion—one chip for every human on Earth (Kelly, 1998: 11). This number is growing rapidly. It will almost certainly approach the hundred billion mark within the decade, and some see the trillion-chip milestone as only a matter of time.[1] Of those already in place, only a tiny fraction (about 3%) live in computers (Kelly, 1998: 11). The rest find homes in an amazing diversity of products: consumer electronics, of course, but also in things like running shoes, bowling balls and birthday candles. The cost of adding limited information-processing ability to an artifact is rapidly approaching zero, and so-called "smart" products are cropping up everywhere.

---

[1] Lucas, P. "The Trillion Node Network." 1999 MAYA Technical Report, #MTR-00001.

"Smart," of course, is a misnomer. Most products with embedded microchips are really quite dumb. The processor in your typical thermostat can compare values for actual and desired temperature, then forward an "on" or "off" command to your furnace, but that's it. Of course, individual neurons are dumb too, but the brains they compose are often not. The real significance of the spread of information devices is not that any one of them is terribly smart, but that more and more are being connected up into networks. Cell phones now "talk" to web-servers, handhelds converse with hard-drives, cars communicate with global positioning satellites, and credit cards trade bits with bank machines. Demand for interoperability is growing, and as designers of information devices work to provide it, they will be laying the foundation for an information system far vaster than the existing World Wide Web. Cyberspace will spill from our desktops and extend tendrils into our kitchens, our cars, and our clothing. It will be many orders of magnitude more complex than any artifact human beings have ever encountered.

What does all this portend? Science writer Robert Wright articulates a question on many minds when he asks "How [are we] to comprehend an age in which …we find ourselves enmeshed in a huge information-processing system, one that seems almost to have a life of its own, and to be leading us headlong into a future that we can't clearly see, yet can't really avoid?" (Wright, 1999).

Emergence & Design

In the late 1960s, William Ross Ashby studied the effects of increasing connectance on the stability of complex systems. The evidence led him to conclude that "All large complex dynamic systems may be expected to show the property of being stable up to a critical level of connectance, and then, as the connectance increases, to go suddenly unstable" (Dyson, 176). Generally speaking, highly interconnected systems can be expected to generate "non-linear" effects—situations where modest inputs spark disproportionate outcomes. Studies like Ashby's raise fascinating questions about the exponential growth of information networks. Might Cyberspace reach the relevant connectivity threshold and "go suddenly unstable"? What would this mean? What kinds of "emergent" or higher-level phenomena will Cyberspace give rise to? Can the system-level behaviors of next generation information networks be understood in advance of their arrival? Does it make any sense to ask whether we can design or engineer emergent properties?

These questions might seem ill-formed. After all, emergent properties are precisely those that cannot be inferred from the behavior of component parts. As George Dyson puts it, "Emergent behavior is that which cannot be predicted through analysis at any level simpler than that of the system as a whole" (1997: 9). When an algorithm is its own shortest description, mathematicians call it *computationally incompressible*; when a physical system defies reductive analysis, it is thought to be

similarly incompressible, and its behaviors are termed "emergent." If there are no shortcuts to figuring out what a system will do—short of actually watching it unfold—then surely the attempt to forecast its emergent properties is a fool's errand.

That "if," though, is a big "if." Phenomena that resist theoretical analysis for centuries will sometimes, and quite suddenly, yield. Facing the bewildering complexity of Tycho Brahe's astronomical data, Kepler could have thrown up his hands, but he persisted and discovered economical descriptions of the planet's (elliptical) orbits. The difficulties of distilling a compact theoretical description of biological growth in the face of nature's diversity might have deterred Darwin, but he persevered, and came up with natural selection. Calling system-level effects "emergent" may mark the computational complexity of the circumstances that generate them, but it should not serve as an excuse to block attempts at deeper comprehension.

Even if the phenomena that interest us were guaranteed to resist the kind of reductive analysis that would confer strong predictive power, we might still profit from a more qualitative understanding of their tendencies. No algorithm will allow us to predict the behavior of a two-year old child, but that does not prevent us from discerning habits and describing personality. Why not strive for similar knowledge of distributed information systems? Seismologists unable to forecast coming earthquakes can nevertheless tell us instructive things about the statistical distribution of earthquakes of various sizes. Perhaps the information dynamics of Cyberspace will someday be similarly characterized.

Our question—and the interests that drive it—are far from academic. Designers and engineers already face challenging questions about the system-level implications of specific design decisions. The emerging field of *information architecture* will grapple with these challenges on at least three fronts. The first is that of *system architecture*- the discipline of designing componentized devices for distributed information systems. The second is the *architecture of user interfaces*—the art of designing tools and conventions that empower the users of information technology. The community of information professionals has yet to achieve clarity on the need for a third branch of information architecture—a field that might be called *information architecture proper.* Here we mean the discipline of designing *information itself*—the actual "currency" of information systems. In each of these three domains, the problem of escalating complexity has raised novel challenges—challenges that we believe complexity theorists will help address.

In a sense, we are inquiring into the possibility of *nonequilibrium engineering*—a discipline devoted to the creation and maintenance of dynamic, energy-consuming, order-generating systems (Kauffman, 1995: 20-1). Until recently, economics was nearly synonymous with the neo-classical economics of diminishing returns—a theoretical framework based on the assumption that markets are fundamentally equilibrium or self-stabilizing systems. The complexity-driven economics of

increasing returns changed all that, and now economists can model nonequilibrium markets (Arthur, 1990; Waldrop, 1992). Traditionally, engineers have designed and built predictable, self-stabilizing artifacts. Suddenly, however, they find themselves building nonequilibrium systems. Can a corresponding expansion of engineering be far behind? We are talking here about the design of systems that will evolve and unfold according to a logic all their own—a possibility that will require an expansion of our understanding of design work. Presumably, nonequilibrium engineers will do things like experiment with the rules that govern local interactions in artificial systems, design or redesign the "currencies" that flow through such systems, and attempt to mold the selective pressures that shape the evolution of information "objects." Where else besides the sciences of complexity can engineers turn for guidance on such matters?

Whether or not nonequilibrium engineering emerges as a recognized discipline, the architects of information systems will continue creating design solutions, and the patterns they embody will spread as they are mimicked by others and replicated mechanically. Some of these patterns will grow entrenched, and become the foundation of the information superstructure of tomorrow. This superstructure will likely have an enormous impact on posterity, for the dynamics of complex systems confer enormous leverage on the creative processes that drive revolutionary ferment. Just as the flowering of life-forms in the Cambrian era determined the basic phyla that dominated the animal kingdom for 540 million years, just so might the design patterns that take hold in these early years of the information revolution determine the basic structural features of human reality for generations to come (Kauffman, 1995).

So the question remains. Can systems science help information technologists understand the global implications of local design decisions iterated many times over? Can our growing knowledge of complex systems be used to shape the unfolding of information systems? Is there any sense in which we can *design* the emergent properties of Cyberspace?

Information Dynamics

Clearly, we lack well-defined methods for tackling such questions. The concepts and techniques of system science will someday gain better purchase, but to date, little progress has been made. My aim in this paper is to explore ways of rendering these questions more tractable.

Suppose we begin with a greatly simplified conception of Cyberspace—a toy model for developing our intuitions. Imagine a network of information-processing "nodes," many of them devices, and some of them people. Information flows through this network, generating various patterns of node activity. By "node activity," we mean a

causal sequence that typically involves: (1) the reception of some number[2] of messages at a node's input terminals, (2) the node's "processing" those messages, in one way or another, and (3) the node's sending out one or more messages from its output terminals.[3] In the case of human beings, node activity can be loosely characterized as a causal sequence involving the reception of perceptual stimuli, some kind of brain or nervous system activity, and some kind of information-generating behavioral response.

The "higher-level" phenomena that emerge in such a system will presumably consist in, or at least be driven by, spreading patterns of node activation. As an example, consider an epidemic caused by a successful computer virus. Specific infections are attributable to individual nodes, but the epidemic must be understood as a property of the system writ large. Similar phenomena rely on individual acts of human will to propagate: chain letters must convince their recipients to replicate and distribute them, for example, and ideologies spread by persuading their hosts to evangelize, persuade, or recruit new believers. Anyone who has studied such phenomena knows that they frequently take on a 'life' of their own—which is one way of saying that they are emergent, driven by imperatives not attributable to any individual.

Such phenomena are the subject-matter of the fledgling discipline of memetics, a field that studies how ideas, norms, habits, attitudes and other *imitable* units of cultural inheritance—"memes" in the jargon of the field—replicate and spread (Dawkins, 1976; Blackmore, 1999). My first general conclusion is that memetics will play an important role in understanding the emergent properties of Cyberspace. (The mathematical modeling techniques of epidemiology are also likely to prove important.) Note also that intellectual historians have been studying spreading patterns of node activity in distributed information systems since at least the early nineteenth century, when the German philosopher G.W.F. Hegel released his classic treatment of ideational dynamics, the *Phenomenology of Spirit*. Archaeologists, religious historians and linguists have also developed fruitful ways of exploring memetic phenomena (Diamond, 1997: 381). Those interested in the emergent properties of Cyberspace would do well to seek guidance from these more mature disciplines.

Are we really suggesting that a postulated science of distributed information systems should look to the "soft" or interpretive disciplines for direction? We are. Explaining spreading patterns of node activation in information systems typically requires reference to the semantic content of the messages passed between nodes. For example, one cannot hope to understand the rise of positivism in the early decades of the twentieth century without understanding the message of positivism. Nor can one

---

[2] "Some number" is meant to include zero, in recognition of cases where relevant node activity is not prompted by any incoming messages.
[3] A theoretical framework for analyzing dataflows in distributed information systems is presented in P. Lucas, *An Ecology of Information Devices*. Still in manuscript.

explain the spread of Newtonian mechanics without reference to the content of Newton's Laws. In information systems, the semantic content of messages ranks among their most important causal properties, and this suggests that the human or interpretive sciences have important contributions to make to a future science of information dynamics.

It is important to understand that distributed information processing systems are not new on the historical scene. In a recently released book and related article, Robert Wright examines how innovations like agriculture, writing and the printing press accelerated the pace at which human cultures process information (2000; 1999). Viewed from this perspective, the internet is just the latest of a long string of developments that is weaving humanity into a tighter and tighter information-processing network—in his words, a "vast social brain" in which we are the neurons (Wright, 1999). Dyson develops a similar theme in a well-argued book provocatively subtitled "The Evolution of Global Intelligence" (1997). Such research underscores the fact that Cyberspace is not an unprecedented phenomenon, but an extension of an increasingly well-understood historical unfolding. Those who wish to understand the emergent properties of Cyberspace would do well to draw from historically informed work on the evolution of precursor information systems.

It is instructive to examine a decisive moment in the development of information systems. When Greek scribes modified the Hebrew *aleph-beth* into the first fully phonetic alphabet, widespread literacy radically changed the rate at which ideas spread and catalyzed one another (Abrams, 1996). Because those ideas could be recorded, copied and distributed, the "space" of ideas gained a kind of independence from concrete material circumstances. The Greek philosophers, who were particularly enamored of the mathematical discoveries of Pythagorus, set out to explore and map this space—a project that culminated in Plato's notion of the purely intelligible realm of Forms. It is illuminating to read the history of western philosophy—once described as 'a series of footnotes to Plato'—as a prolonged effort to gain clarity about the nature and properties of knowledge or information systems—of "Cyberspace," if you will. Unfortunately, this effort has been hampered by the assumption that systems of knowledge are, if not immutable, then at least squarely within what complexity theorists would call the "ordered regime"—that portion of a system's state-space in which higher-level patterns can be counted on to settle down, die out, or remain predictably linear. (Plato, Descartes and Kant are prime examples here, but by no means the only culprits.) Before Hegel, the possibility that knowledge might form a dynamic, nonequilibrium system received little attention. Unfortunately, this possibility still awaits adequate philosophical treatment.

In the meantime, intellectual historians have compiled a wealth of data on the dynamics of information systems. Consider, for example, Thomas Kuhn's work on the rise and fall of scientific paradigms. Complexes of concepts, presuppositions and experimental practices define research traditions that remain stable for a time—periods Kuhn called "normal science"—until they are overthrown by competing paradigms in

periods of disciplinary transformation Kuhn called "revolutions." This "punctuated equilibrium" structure suggests that some information systems (sciences in particular) reside in a state of self-organized criticality, evolving and building complexity in much the way biological systems do.

When it comes to cyberspace, this conjecture can be tested. Imagine a sensor capable of monitoring the quantity of data traffic through central internet servers or routers. Per Bak (1996) has argued that systems in a state of self-organized criticality generate a distinctive signal-pattern known as '1/f noise.' This means that if the frequency of the various data-loads turns out to vary inversely with the size of those loads in accordance with some kind of power law, we would have strong reason for thinking that Cyberspace does indeed reside in a state of self-organized criticality.

Information Architecture & Non-Equilibrium Design

Let us now approach the problem from the other end. How do those who work on the cutting edge of information architecture conceptualize the problem of designing for a networked world? What sorts of "high-level" properties do they want information systems to have, and how do they approach the task of designing them? What follows is a (necessarily) brief status report from the front lines of information architecture.

Moore's Law records a doubling of processor speed every eighteen months for the last quarter century. Every couple of years, there is a comparable doubling of the size and complexity of embedded software systems (Ommering, et al, 2000). These trends have made it increasingly hard to engineer *tractable* information systems. The millions of lines of code that make up the Windows operating system create a software environment that is so mind-bogglingly complex that even the best-designed applications cause system crashes. The capacity of human software engineers to manage such complexity is nearing its limits, and pressure is mounting to institute a more rigorous approach to managing engineered complexity.

The solution is not hard to describe—not in general terms, anyway. Engineers have long understood that the effective or operational complexity of a device can be reduced through componentization: the designer's job becomes an order of magnitude easier if she can plug in components and count on them to meet precisely defined functional specifications. However, because software is so cheap to write, the gospel of componentization has been slow to spread. Thankfully, movements such as object-oriented programming signal that the world of software engineering is assimilating the idea of rigorous componentization, and reducing the operational complexity that plagues the process of software design. As a result, we should see improvement in the tractability of complex software systems.

What other high-level properties are systems designers seeking to engineer? Among them are two that componentization also serves to promote: adaptability and evolvability. Information systems that cannot adapt to the pace of change that they themselves make possible are destined to be replaced by systems that can, and it is much simpler to upgrade or adapt a rigorously componentized system. As engineers work out more effective ways to componentize complex devices, they are in fact designing systems with the global properties of adaptability and evolvability. Paradoxically, the quest for tractability leads engineers to create systems that are more and more adaptive, more and more life-like, and harder and harder to micromanage—a process Kevin Kelly documents in his provocative book *Out of Control* (1994).

This paradox hints at an important tension. The self-organization that complexity theorists study and strive to simulate is *exactly* the kind of thing that engineers have traditionally sought to eliminate from designed systems. Per Bak writes of the "seeming intractability of emergent phenomena," and engineers testify to the difficulty of designing self-configuring components that perform in accord with functional specifications (1996: 7). Until very recently, engineering was *precisely* the art of creating systems that would *not* give rise to self-organization or other intractable phenomena. In a sense, this means that engineers and complexity theorists have been working the same problem from opposite ends. Each has been exploring an important boundary in an abstract state-space for artificial systems: the place where tractability merges into intractability, where emergence emerges. While engineers have approached self-organization as something to be designed *against*, complexity theorists have treated it as something to be designed *for*. Each should be able to learn from the successes and failures of the other.

Let us turn now to the user interface front, where designers are challenged to limit the effective complexity that confronts the users of information devices. Software designers are under constant pressure to integrate new features, but the capacity of users to sort through the visual clues that crowd an interface has not changed radically. Most of us have experienced the frustration of having to hunt through dozens of pull-down menus for the feature we need, and in each of our cases, there is a limit to how much of this we are willing to put up with. To empower the users of information technology, interface designers must find creative ways to make their tools transparent extensions of the user. Good tools *background* themselves, allowing users to focus on the task at hand. As in the case of system architecture, good design here involves reducing effective complexity: decomposing complex experiences into easily navigable arrangements of simpler ones.

What we have called information architecture proper—the practice of designing information itself—has yet to emerge as a recognized discipline. As information devices are networked, however, data is increasingly "freed" from particular devices: it makes less and less sense to think of information as being "in" my computer or "in" the electric meter, and more and more sense to think of the devices as "in" the

information.[4] Information "lives" in Cyberspace, and particular devices are incidental to its existence (just as pure ideas belonged to Plato's intelligible realm, their material manifestation being incidental to their existence). As we grow increasingly comfortable with this way of thinking—and it *will* happen—we will gradually come to see information as something that needs designing. Information architecture proper will come into its own, and we will design information "objects" that allow us to see the world in new and fascinating ways.

Where does this leave us? With more questions than answers, no doubt. Still, one must ask the important questions before one can answer them. In the quote that heads this paper, George Dyson notes that "the proliferation of microprocessors and the growth of distributed communications networks hold mysteries as deep as the origins of life [and] the source of our own intelligence." We contend that these same phenomena raise equally deep *practical* questions. As information professionals learn to design for a radically networked world, they will stretch the boundaries of their discipline, and create the field of non-equilibrium design. By helping them understand how local interactions give rise to system-level behavior, complexity theorists will have a hand in this creation.

---

[4] Lucas, P. "Pervasive Information Access and the Rise of Human-Information Interaction." Paper presented April 5 at CHI 2000, held at The Hague, Netherlands.

Information Architecture and the Emergent Properties of Cyberspace

Works Cited

Arthur, Brian. (1990) "Positive Feedbacks in the Economy." *Scientific American* (February, 1990): 92-99.

Bak, Per. (1996) *How Nature Works*. New York, Copernicus.

Blackmore, Susan. (1999) *The Meme Machine*. Oxford University Press.

Dawkins, Richard. (1976) *The Selfish Gene*. Oxford University Press.

Diamond, Jared. (1997) *Guns, Germs, and Steel*. New York, W.W. Norton & Co.

Dyson, George. (1997) *Darwin Among the Machines*. Reading, Massachusetts, Perseus Books.

Kelly, Kevin. (1994) *Out of Control*. Reading, Massachusetts, Perseus Books.

----- (1998) *New Rules for the New Economy*. Middlesex, England, Viking Penguin.

Lucas, Peter. *An Ecology of Information Devices*. (in manuscript)

Lucas, Peter. "Pervasive Information Access and the Rise of Human-Information Interaction." Paper presented April 5 at CHI 2000, held at The Hague, Netherlands.

Lucas, Peter. "The Trillion Node Network." 1999 *MAYA Technical Report*, #MTR-00001.

Ommering, R., van der Linden, F., Kramer, J., and Magee, J. (2000) "The Koala Component Model for Consumer Electronics Software." *Computer*. pp. 78-85.

Waldrop, Mitchell. (1992) *Complexity.* New York, Simon & Schuster.

Wright, Robert. (2000) *Non-Zero: The Logic of Human Destiny*. New York, Pantheon Books.

----- "The Web We Weave." *Time*. Dec. 31, 1999