



Everything, the universe, and life

by N. Gershenfeld

The previous MIT Media Laboratory special issue¹ of the *IBM Systems Journal* closed with a group of papers describing an emerging physical science effort in the Media Lab. These inverted the historical focus of the Media Lab, which was on the revolutionary implications of a digital representation for freeing the content of information from the constraints of its physical representation. When the Media Lab was founded, the debate over the frame size and rate for the next video standard was widely seen as a matter of great significance; the Lab argued instead that a scalable encoding would allow a television set to be downloaded with a TV signal to match the needs of the signal and its viewer. This once-radical idea is now ubiquitous in streaming digital media.

In that issue's concluding essay² I argued that the very success of this agenda presents an even-grander challenge of merging the best features of the digital worlds we are creating with those of the physical world we are born into. While billions of dollars have been spent on ever-faster CPUs, these CPUs are put into cases that have changed little from the earliest days of computing. For more and more people, and problems, the greatest constraint in computing is getting the right information to the right place at the right time, rather than the speed with which it can be processed. In retrospect, the return to consideration of the mechanisms for physical interaction with information can be understood as just the ultimate expression of the original meaning of "media." The accompanying papers in that issue offered preliminary glimpses of how to bring together newly freed bits with more-capable atoms by thinking outside of the computer box—ranging from a calculation of the thermodynamic implications of the entropy associated with the information in the electronic states of

a string of bits,³ to a description of how those bits could be carried through the ionic conductivity of the human body in order to attach digital content to physical gestures.⁴

The corresponding papers in the current issue show that the integration of bits and atoms has matured from an elusive metaphor to an explicit research agenda, over length scales ranging from 10 to 10^{-10} meters. Maguire et al. discuss the means to manipulate coherent quantum information stored in atomic nuclei; Post et al. weave information into the fabric of clothing with process technology using conducting threads for integrated circuit interconnect; Paradiso et al. describe active architectural surfaces; and Omojola et al. present an entire gallery in the Museum of Modern Art (MoMA) as a responsive user interface.

Along with this range of length scales comes a range in the scale of system complexity. Where Maguire et al. are concerned with the fundamental question of the representation of a bit of information, Post et al. study the connection of a few processing elements, and Omojola et al. describe a table containing 400 microcontrollers. These scales are mirrored in levels of abstraction; the MoMA installation comprised languages and programs at seven levels of description: software radio-frequency instrumentation, local data concentration, network communications transport, real-time data analysis, object encapsulation, media management, and graphical and interaction design.

©Copyright 2000 by International Business Machines Corporation.

Each of these projects provides an example of the successful integration of the information in a system with its physical properties, across an enormous range of scales. The problem with them is the extent to which they all satisfied their initial design goals. A hint as to why this might be a cause for serious concern appears toward the end of the paper by Omojola et al., where the authors explain that the hardware development for the installation was a relatively straightforward process, but that the software debugging proceeded right up to the show's opening (and beyond). Having problems with the testing of system software is certainly not a recent development; what is new is the prospect that these problems will be magnified as systems grow from thousands to millions to billions of components. The ways we now manage networks of interacting elements do not scale to a world in which every shoe and sleeve can contain a communicating computer.

There are a number of existing examples of large-scale distributed systems, most notably the telephone network, the power grid, and the Internet. Each of these already shows rather disturbing emergent failure modes.⁵ On January 15, 1990, AT&T's long-distance network went down for 9 hours, leaving an estimated 65 million calls uncompleted. The problem was a system-wide software upgrade that was intended to improve the reliability of the 4ESS digital circuit switch. The upgrade eliminated the need for a switch resuming service after a failure to explicitly notify its neighboring switches of the change in its state, since that is already conveyed by the first new traffic that the switch generates. The failure started when a switch in New York City suffered a minor hardware glitch, which caused it to go off line. After a brief transient it returned to service, and started passing messages to adjacent switches. When the first new message arrived, these neighbors began executing the upgraded routine to update their status record for the switch that had been down. But, when a second message came from the formerly faulty switch before the error routine had finished executing, a bug in that routine caused the neighboring switches to crash. These then returned to service, tripping their neighbors when they did so, thereby initiating a chain reaction that propagated through the entire phone system. Every switch was doomed to repeat the failure of its neighbors.

Something similar happened to the electrical grid. On August 10, 1996, a power line in Oregon that was sagging due to heat grounded through a tree. Because a few other backup lines were already down,

this produced a large voltage transient and triggered automatic load-shedding. This loss of service increased the demand on adjacent power lines, which in turn were automatically shut down by software intended to protect them from transients. The resulting instability propagated through the grid's Pacific Intertie, eventually knocking out 25 gigawatts of nuclear and conventional generating capacity, in 190 plants serving 7.5 million customers over 9 states.

The largest distributed system of all, the Internet, has seen numerous comparable failures. A dramatic example happened on April 23, 1997, when most domestic Internet service providers lost contact with the network's backbone, shutting down all their traffic. A small backbone provider, MAI Network Services in McLean, Virginia, accidentally communicated incorrect routing tables to the network providers through the Border Gateway Protocol. Because these were missing Autonomous Systems Path identifiers, they had the effect of identifying MAI as the shortest path to 25 percent of the Internet for all users. Even worse, because the routing information was marked as being authoritative, it was trusted over existing routes. The erroneous tables automatically propagated between routers, eventually turning 40 percent of the network into a "black hole" that swamped MAI with lost packets. It took hours to chase down and replace the incorrect information being passed around the network.

In each of these examples, the failure mode was an emergent property of the entire system. The individual elements were sufficiently intelligent to be able to take action to correct local errors, and sufficiently well connected for those to cascade into a global fault. What was notably missing was any ability for the system components to recognize that they were slavishly following flawed instructions, so that life experience along with erroneous information could travel between them.

Return now to the preceding series of papers in this issue. They provide a roadmap for embedding networked communications in every thing around us—down to its atomic structure—and insight into the essential implications of such a world for interacting with information, as described by Ullmer and Ishii, and Selker and Burleson. But they provide no guidance at all as to how to manage such systems. To the contrary, the telephone, power grid, and Internet examples make clear the extent to which complex systems based on present engineering principles can become unmanageable.

The answer is to recognize that function as well as failure must become an emergent property of these interacting systems. The problem with papers describing devices that do meet design specifications is that the devices are not given the means to transcend the limitations of those advance specifications. Scalable design for systems of billions of things must include the ultimate consequence of merging bits and atoms: lifelike attributes. As Pentland's introduction suggests, the individual elements must be able to adapt to changing environments, modify their behavior autonomously, evolve operational instructions from initial conditions rather than a complete description, and so forth.

These are familiar characteristics of living systems, but not yet matters of engineering design. While the integration of fine-grained digital systems with rich sensory interfaces can help provide insight into the functioning of biological systems,⁶ it is the insight into biology that is needed to scale the integration of bits and atoms. Preliminary examples do suggest that basic scaling laws can be turned into engineering design principles that allow the robust operation of a system to be an emergent property of unreliable components.⁷ The challenge for the next special issue of the *IBM Systems Journal* is to develop an engineering theory of applied life.

ture of Mathematical Modeling, and *The Physics of Information Technology*. Dr. Gershenfeld has a B.A. degree in physics with High Honors from Swarthmore College, a Ph.D. from Cornell University, was a Junior Fellow of the Harvard University Society of Fellows, and a member of the research staff at AT&T Bell Laboratories.

Cited references

1. *IBM Systems Journal* **35**, Nos. 3&4 (1996).
2. N. Gershenfeld, "Physics and Media," *IBM Systems Journal* **35**, Nos. 3&4, 575–577 (1996).
3. N. Gershenfeld, "Signal Entropy and the Thermodynamics of Computation," *IBM Systems Journal* **35**, Nos. 3&4, 577–586 (1996).
4. T. G. Zimmerman, "Personal Area Networks: Near-Field Intra-body Communication," *IBM Systems Journal* **35**, Nos. 3&4, 609–617 (1996).
5. P. G. Neumann, "Computer-Related Risks," Addison-Wesley Publishing Co., Reading, MA (1995).
6. S. R. Manalis, E. B. Cooper, P. F. Indermuhle, P. Kern, P. Wagner, D. G. Hafeman, S. C. Minne, and C. F. Quate, "Microvolume Field-Effect pH Sensor for the Scanning Probe Microscope," *Applied Physics Letters* **76**, 1072–1074 (2000).
7. J. R. Heath, P. J. Kuekes, G. S. Snider, and R. S. Williams, "A Defect-Tolerant Computer Architecture: Opportunities for Nanotechnology," *Science* **280**, 1716–1721 (1998).

Neil Gershenfeld *MIT Media Laboratory, 20 Ames Street, Cambridge, Massachusetts 02139-4307*. Dr. Gershenfeld leads the Physics and Media Group at the MIT Media Laboratory and directs the Things That Think research consortium. His laboratory investigates the relationship between the content of information and its physical representation, from developing molecular quantum computers, to smart furniture, to virtuosic musical instruments. Author of the books *When Things Start to Think*, *The Na-*