wave sensors and radars); cloud cover and cloud phase (from traditional passive visible and infrared measurements and from newer active instruments); snow cover, sea ice cover, and soil moisture (mostly from passive microwave instruments, which can sense through clouds); smoke, dust, volcanic ash, and other aerosols (from a variety of passive sensors); and gases, including carbon dioxide, sulfur dioxide, and ozone (from infrared and ultraviolet spectrometers). All of these quantities will gradually become part of improved NWP models and forecasts.

In the 50 years since TIROS 1, many fields other than meteorology have benefited from space-based observations, including atmospheric science, climate studies, oceanography, hydrology, ecology, and geology (*9*). It is becoming increasingly clear that these fields are interrelated. When a volcano erupts, the dust becomes a meteorological problem and a threat to aviation. Changes in sea surface temperatures alter the tracks of storm systems. Changes in atmospheric composition affect climate and air quality. Future satellite sensors will thus serve a variety of fields. A list of missions has been recommended as a result of the U.S. National Research Council's decadal survey for Earth science (*10*, *11*), and many other missions are planned by other countries (*12*, *13*).

The first 50 years of space-based Earth observation progressed from crude observations to scientific understanding to stewardship of the atmosphere and of Earth. Today's space-based observations will likely appear crude in 50 years (*14*). The new observations will result in many scientific insights and should help humanity to weather what could be the worst of global warming and other environmental problems.

References and Notes

1. J. G. Vaeth, *Weather Eyes in the Sky: America's Meteorological Satellites* (Ronald Press, New York, 1965).
2. H. Wexler, *J. Br. Interplanet. Soc.* **13**, 269 (1954).
3. P. K. Rao, S. J. Holmes, R. K. Anderson, J. S. Winston, P. E. Lehr, *Weather Satellites: Systems, Data, and Environmental Applications* (American Meteorological Society, Boston, 1990).
4. M. J. Bader *et al.*, Eds., *Images in Weather Forecasting: A Practical Guide for Interpreting Satellite and Radar Imagery* (Cambridge Univ. Press, Cambridge, UK, 1995).
5. R. C. Sheets, *Weather Forecast.* **5**, 185 (1990).
6. S. Q. Kidder, T. H. Vonder Haar, *Satellite Meteorology: An Introduction* (Academic Press, San Diego, CA, 1995).
7. Monthly statistics on the data used in U.S. NWP models are at www.nco.ncep.noaa.gov/sib/counts.
8. A. J. Simmons, A. Hollingsworth, *Q. J. R. Meteorol. Soc.* **128**, 647 (2002).
9. National Research Council, *Earth Observations from Space: The First 50 Years of Scientific Achievements* (National Academies Press, Washington, DC, 2008).
10. National Research Council, *Earth Science and Applications from Space: National Imperatives for the Next Decade and Beyond* (National Academies Press, Washington, DC, 2007).
11. National Research Council, *Satellite Observations to Benefit Science and Society: Recommended Missions for the Next Decade* (National Academies Press, Washington, DC, 2008).
12. World Meteorological Organization (WMO), www.wmo.int/pages/index_en.html.
13. Committee on Space Research (COSPAR), http://cosparhq.cnes.fr.
14. National Research Council, *America's Future in Space: Aligning the Civil Space Program with National Needs* (National Academies Press, Washington, DC, 2009).

10.1126/science.1185867

ENGINEERING

Intelligent Infrastructure for Energy Efficiency

A substantial fraction of wasted energy can be recovered by extending insights from the architecture of the Internet to the infrastructure of buildings.

Neil Gershenfeld,[1] Stephen Samouhos,[1,2] Bruce Nordman[3]

Buildings use 40% of the primary energy supplied in the United States, and more than 70% of all generated electricity (*1*), primarily for heating, cooling, and lighting. About 20% of the energy used by buildings can potentially be saved by correcting faults, including malfunctions and unnecessary operation (*2*). Initial deployments of advanced control systems currently in development suggest that they can save an additional 10 to 20% (*3*). The energy efficiency resource recoverable through such improved building controls and fault detection corresponds to the output from hundreds of power plants, equivalent to more than one-third of the coal-fired power production in the United States (*1*). Realizing these substantial savings will require introducing intelligence into the infrastructure of buildings, to distribute the optimization of their operation and detection of their faults.

Intelligent infrastructure extends "smart grid" initiatives that seek to save energy by allowing utilities to manage loads, such as turning off air conditioners during peak demand (*4*). However, a grid cannot be smart if it is connected to dumb devices. Currently, modifying a building is costly and labor-intensive; it can cost $1000 to add a control point containing a $1 sensor to a building, requiring a skilled installer to connect it to a central controller that then must be reconfigured. This situation is analogous to computing and communications before the Internet, when terminals and telephones were connected to mainframes and central office switches.

The Internet allowed applications to reside where information is created and consumed, from reading e-mail to viewing virtual worlds. In this way, its applications are independent of how the network connecting them is constructed (*5*). This observation is equally applicable to building infrastructure: Sensors and actuators can compute and communicate to solve problems locally rather than having functions fixed by a central controller (*6*).

Many of the candidate standards for smart building systems are recreating rather than extending the development of the Internet. There are, however, important differences between high-performance buildings and networks. Installation lifetimes for buildings are measured in decades, and the cost of installation and maintenance can dwarf the cost of devices. Events can happen over seasons (such as cooling versus heating) rather than seconds, requiring efficient handling of slow rather than fast events. Operation must often be unsupervised, air and water as well as information must be moved, and people are an integral part of the building system. Each of these practical considerations presents new research challenges; they cannot be addressed simply by making better use of available technologies.

A testbed for intelligent infrastructure for energy efficiency (I2E), shown in panel A of the figure, consists of 100 nodes (*7*) that cost about $1 each but contain interfaces for sensors and loads, implement Internet Protocol (IPv6) communications over the dc control wiring already in the building, and provide an

[1]Center for Bits and Atoms, Massachusetts Institute of Technology, Cambridge, MA 02138, USA. [2]Building Technologies, Massachusetts Institute of Technology, Cambridge, MA 02138, USA. [3]Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA. E-mail: gersh@cba.mit.edu, bnordman@lbl.gov, stratos@mit.edu

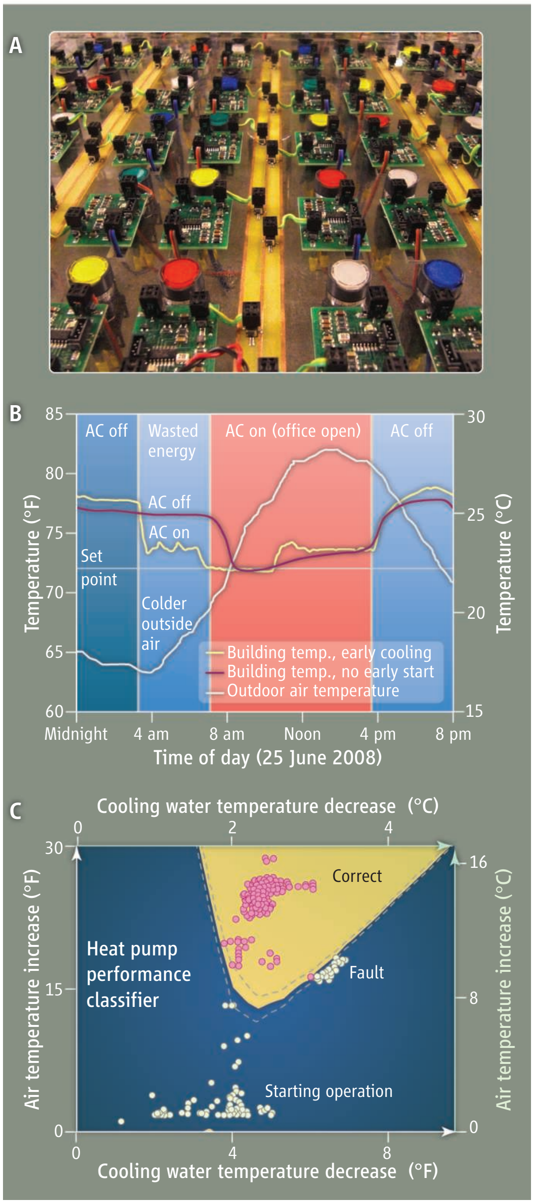

Bringing energy loss under control. (**A**) A key element for intelligent infrastructure for energy efficiency (I2E) is sensing and controlling the state of a building. The testbed shown has 100 embedded IPv6 Web servers, each with sensor inputs and control outputs, at a component cost of roughly $1. (**B**) Measured responses, based on I2E testbed data, for an office building with an oversized air conditioner that turns on 4 hours early (blue), versus an appropriate start time closer to the start of the workday (red). (**C**) Classification of a heat pump's operation based on analysis of drops in water temperature and increase in air temperature. Pattern recognition with a convex kernel method separates clusters of normal operation (yellow region) from starting operation conditions, as well as a separate set of responses that appear to reflect faulty operation (blue region).

embedded Web server. At wiring junctions, changes in impedance will cause part of a signal to be reflected, which is prevented in high-speed networks by adding active terminations. However, signal impulses in building wiring will decay on the order of the time it takes them to transit the building, typically microseconds. By communicating to embedded devices with transient impulses slower than this rate (megabits per second), a network can be connected in an arbitrary way (*8*). This approach minimizes complexity rather than maximizes capacity and supports any available signaling medium and data rate. A key feature of the scalability of the Internet—its ability to operate in the same way as more connections and devices are added—has been its use of protocol definitions (*9*) that do not limit performance numbers, unlike candidate standards for building systems that fix these parameters (*10*).

Another aspect of scalability has been hierarchical routing, which can provide global connectivity even for low-power devices with unreliable connections (*11*). By embedding Web servers in individual sensors and actuators to present and control their state, a new device can document its own capabilities, rather than have its functions defined in advance as done today (*12*).

An example of how this works is shown in panel B of the figure, which presents data collected from sensors in an I2E testbed installation, an office building at MIT. Air conditioners with excess capacity for the actual building load were automatically turning on at 4:30 a.m. to cool the building for a workday starting 4 hours later, even though the building could be cooled in 30 min. Even worse, in the early morning the outside air was colder than the desired temperature; the most efficient strategy would have been to bring in outside air. Efficient building air handlers do have dampers to regulate the mixing of outside air; however, another I2E testbed study found a mixing unit failing to mix because of unexpected internal recirculation.

To be able to respond to such changing conditions, an I2E system must integrate control of all of the degrees of freedom in a building and adapt to their actual configuration. If networked intelligence is embedded into all of an I2E system's devices, the system can use them to implement controllers that have the same structure as the system being controlled (*13*).

Historically, building infrastructure faults have been detected (if at all) with simple setpoint alarms; more recently, machine learning techniques have been applied that require supervised training (*14*). Modern methods for pattern recognition, such as convex kernel methods (*15*), allow anomalies to be found in unlabeled data without iterative learning procedures; panel C of the figure shows such a decision surface found for a heat pump that identifies its faulty operation.

There are also opportunities for saving energy with dense networks of intelligent actuators as well as sensors. Heating and cooling pumps and fans have been found to contribute more than 50% of the HVAC-related energy usage in commercial buildings (*16*). This includes kinetic energy in moving air and water, drag in pipes and ducts, and internal losses. Efficiency is maximized by minimizing velocity and drag, which can be accomplished by distributing a greater number of smaller impellers over internal surfaces, rather than just a few large ones at inlets or exhausts (*17*).

The energy benefits in these I2E systems will be negated if their control systems consume more energy than they save. Because the relevant operational time scales are measured in hertz rather than gigahertz, most of the clock cycles in a conventional processor will be wasted. This static power consumption can be minimized by using asynchronous logic that is driven by events rather than by a clock. Because building architecture changes much more slowly than computer architecture, the embedded computing should ideally be reconfigurable as well as asynchronous, so that choices such as word sizes or instruction sets are not fixed in the construction of a building. These constraints suggest the use of reconfigurable asynchronous logic in building infrastructure, an area of active research originally developed for use in high-performance computing (*18*).

I2E faces operational as well as technological hurdles to widespread adoption. Building codes evolve slowly to reflect worst-case experiences and are applied on a case-by-case basis, whereas computer codes are developed much more rapidly and are debugged by finding and fixing errors. Construction costs typically do not reflect life-cycle operational costs, which are accounted for independently. Individual building controls do not currently have access to economic data that could convey to occupants the costs of their actions (in dollars, CO_2 emissions, or utility load).

Bridging between these worlds will not be accomplished with an investment at a single point in the system. Intelligent building infra-

Published by AAAS

structure can be thought of as a long "green" tail distribution of many small savings that add up to a major opportunity for reducing energy consumption, while also improving a building's responsiveness to its occupants. This presents corresponding challenges that are at the frontiers of distributed computing and communications; rather than replicating the history of their development, today's best practices can be extended to this largest of all programming environments—the built environment.

References and Notes

1. U.S. Department of Energy, *Buildings Energy Data Book* (http://buildingsdatabook.eere.energy.gov, 2009).
2. K. W. Roth *et al., Energy Impact of Commercial Building Controls and Performance Diagnostics*, TIAX (http://apps1.eere.energy.gov/buildings/publications/pdfs/corporate/pnnl-15149_market_assessment.pdf, 2005).
3. M. R. Brambley *et al., Advanced Sensors and Controls for Building Applications: Market Assessment and Potential R&D Pathways* (http://apps1.eere.energy.gov/buildings/publications/pdfs/corporate/pnnl-15149_market_assessment.pdf, 2005).
4. National Institute of Standards and Technology, Smart Grid Interoperability Standards Project (http://nist.gov/smartgrid/).
5. J. H. Saltzer *et al., ACM Trans. Comput. Syst.* **2**, 277 (1984).
6. N. Gershenfeld *et al., Sci. Am.* **291**, 76 (2004).
7. D. Kopp, Sarvanet, personal communication (2009).
8. N. Gershenfeld, D. Cohen, *IEEE Circ. Dev.* **22**, 48 (2006).
9. DARPA Internet Program Protocol Specification RFC 791 (www.ietf.org/rfc/rfc0791.txt, 1981).
10. ZigBee Alliance (www.zigbee.org/).
11. Internet Engineering Task Force (http://tools.ietf.org/html/draft-ietf-roll-of0-00).
12. BACnet (www.bacnet.org/).
13. C. Langbort *et al., IEEE Trans. Automat. Contr.* **49**, 1502 (2004).
14. M. Krarti, *J. Sol. Energy Eng.* **125**, 331 (2003).
15. A. Rahimi, B. Recht, in Proceedings of the 8th European Conference on Computer Vision Workshop on Statistical Learning in Computer Vision, Prague, 2004 (http://groups.csail.mit.edu/vision/vip/papers/rahimi-ncut.pdf).
16. California Measurement Advisory Council (www.calmac.org, 2003).
17. A. Sun, unpublished results.
18. N. Gershenfeld *et al.*, in POPL '10 (ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages), 17 to 23 January 2010, Madrid (http://phm.cba.mit.edu/papers/09.11.POPL.pdf).

10.1126/science.1174082

MOLECULAR BIOLOGY

Reliable Noise

David Levens[1] and Ashutosh Gupta[1,2]

Assessing how the noise created in transcription factor regulatory circuits affects gene expression is essential to understanding network operation and output.

Most measurements of gene expression assess large numbers of cells to improve precision and reduce the "standard error" (the standard deviation of the mean). Yet, the standard deviation of the fluctuations of a measured property, such as cell proliferation, over time in a single cell (in a system at equilibrium or steady state) or across a cell population, scaled to the mean of the measured property, is defined as "noise." Despite this pejorative, a full accounting of noise provides insights into the pathways and mechanisms controlling a measured property. On page 1142 of this issue, To and Maheshri demonstrate that noise itself can generate a system that switches spontaneously between high and low gene expression (*1*). This finding implies that fluctuation in the numbers of regulatory molecules may drive physiological transitions without having to precisely specify the numbers of other molecules needed to prepare chromatin and make RNA. However, these same fluctuations might initiate and sustain pathological states, so mechanisms to suppress such fluctuations must also exist.

The basic experimental scheme used by To and Maheshri involves expressing TetVP16, a recombinant transcription factor, from a weak minimal promoter bearing either one or seven binding sites for TetVP16 itself. This positive feedback arrangement mimics a commonly occurring biological regulatory motif (*2*). Reporter genes encoding fluorescent proteins that are driven by either promoter are then used to monitor transcriptional output in cells. Upon the graded removal of doxycycline, a compound that inhibits the binding of TetVP16 to DNA, a cell population can transition from low to high reporter gene expression in this system.

To and Maheshri observed that with a single TetVP16 binding site in the promoter, the entire cell population increased reporter gene expression gradually and coherently. However, with seven TetVP16 binding sites, even at low doxycycline concentrations, a single burst of transcription could drive enough TetVP16 expression to enable visualization by sustained high reporter expression, the result of a high transcription output state in individual cells that sporadically relaxed to low output. The high-output state was associated with bursts of transcription that were less frequent, and either of longer duration or of higher intensity compared to bursts observed in the low-output state. Because only a single polymerase can initiate transcription at a promoter at one time, prolonging a burst to include more successive rounds of transcription initiation will increase output relative to more frequent, but short bursts.

The same bimodal pattern of low and high gene expression was also observed with a stronger promoter bearing just one TetVP16 binding site. Thus, cooperative binding of TetVP16 was excluded as the cause of the switch from low to high transcription output. This two-state system is distinct from the monotonic curve defined by calculations based solely upon the binding constants and concentrations of interacting molecular species in the absence of cooperativity. The study of To and Maheshri also reveals the difficulty of rigorously accounting for the biologically relevant species of macromolecules. To construct a working mathematical model, the authors had to correct for a cytoplasmic reservoir of inactive TetVP16 molecules and for

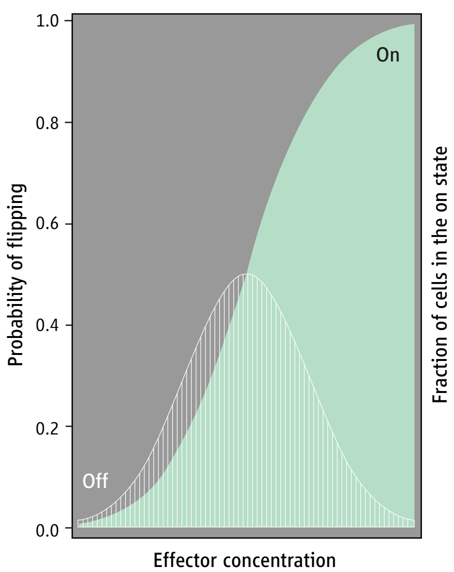

Dynamic system stability. The fracton of cells expressing a gene is a function of the concentration of an effector molecule (as in the system used by To and Maheshri). At very low or high effector concentrations, the expression system is often off (gray) or on (green). At intermediate concentrations, the system is bimodal, flipping between both states (region with lines). Nevertheless, the overall system is stable.

[1]Laboratory of Pathology, Center for Cancer Research, National Cancer Institute, Bethesda, MD 20892, USA. [2]Department of Physics, University of Maryland, College Park, MD 20742, USA. E-mail: levensd@mail.nih.gov