Reverse Engineering of Biological Complexity
Marie E. Csete, *et al.*
Science **295**, 1664 (2002);
DOI: 10.1126/science.1069981

readable model definition that enables models to be exchanged between software tools. Systems Biology Workbench (SBW) is built on SBML and provides a framework of modular open-source software for systems biology research. Both SBML and SBW are collective efforts of a number of research institutions sharing the same vision (25).

How does the idea of systems biology impact pharmaceutical industries and medical practice? The most feasible application of systems biology research is to create a detailed model of cell regulation, focused on particular signal-transduction cascades and molecules to provide system-level insights into mechanism-based drug discovery (26–28). Such models may help to identify feedback mechanisms that offset the effects of drugs and predict systemic side effects. It may even be possible to use a multiple drug system to guide the state of malfunctioning cells to the desired state with minimal side effects. Such a systemic response cannot be rationally predicted without a model of intracellular biochemical and genetic interactions. It is not inconceivable that the U.S. Food and Drug Administration may one day mandate simulation-based screening of thera-

peutic agents, just as plans for all highrise building are required to undergo structural dynamics analysis to confirm earthquake resistance.

Although systems biology is in its infancy, its potential benefits are enormous in both scientific and practical terms. A transition is occurring in biology from the molecular level to the system level that promises to revolutionize our understanding of complex biological regulatory systems and to provide major new opportunities for practical application of such knowledge.

References and Notes

1. N. Weiner, *Cybernetics or Control and Communication in the Animal and the Machine* (MIT Press, Cambridge, MA, 1948).
2. H. Kitano, *Foundations of Systems Biology* (MIT Press, Cambridge, MA, 2001).
3. E. H. Davidson et al., *Science* **295**, 1669 (2002).
4. Examples of such databases are: Signal Transduction Knowledge Environment (STKE http://www.stke.org/); KEGG (http://www.genome.ad.jp/); EcoCyc (http://ecocyc.org/).
5. M. Eisen et al., *Proc. Natl. Acad. Sci. U.S.A.* **95**, 14863 (1998).
6. S. Chu et al., *Science* **282**, 699 (2000).
7. S. Onami et al., in *Foundations of Systems Biology*, H. Kitano, Ed. (MIT Press, Cambridge, MA, 2001), pp. 59–75.
8. S. Imoto et al., *Pacific Symposium on Biocomputing 2002* (World Scientific, Singapore, 2002), pp. 175–186.
9. C. Yoo et al., *Pacific Symposium on Biocomputing 2002* (World Scientific, Singapore, 2002), pp. 498–509.
10. T. Ideker et al., *Pacific Symposium on Biocomputing* (World Scientific, Singapore, 2000), pp. 302–313.
11. J. Edwards et al., *Nature Biotechnol.* **19**, 125 (2001).
12. M. Borisuk et al., *J. Theor. Biol.* **195**, 69 (1998).
13. K. Chen et al., *Mol. Biol. Cell* **11**, 369 (2000).
14. M. E. Csete, J. C. Doyle, *Science* **295**, 1664 (2002).
15. N. Barkai et al., *Nature* **387**, 913 (1997).
16. U. Alon et al., *Nature* **397**, 168 (1999).
17. T.-M. Yi et al., *Proc. Natl. Acad. Sci. U.S.A.* **97**, 4649 (2000).
18. V. Dassaw et al., *Nature* **406**, 188 (2000).
19. G. Weng et al., *Science* **284**, 92 (1999).
20. Alliance for Cellular Signaling (http://www.cellular-signaling.org/).
21. M. Burns et al., *Science* **282**, 484 (1998).
22. R. Anderson et al., *Nucleic Acids Res.* **28**, e60 (2000).
23. P. Simpson et al., *Proc. Natl. Acad. Sci. U.S.A.* **95**, 2256 (1998).
24. P. Gilles et al., *Nature Biotechnol.* **17**, 365 (1999).
25. Additional information can be obtained at http://www.cds.caltech.edu/erato/ or http://www.sbml.org.
26. J. Gibbs, *Science* **287**, 1969 (2000).
27. C. Sander, *Science* **287**, 1977 (2000).
28. D. Noble, *Science* **295**, 1678 (2002).
29. I thank J. Doyle, M. Simon, and members of ERATO Kitano project for fruitful discussions. Supported by the ERATO and BIRD program of the Japan Science and Technology Corporation, and the Rice Genome and Simulation Project of the Ministry of Agriculture, Japan.

REVIEW

Reverse Engineering of Biological Complexity

Marie E. Csete[1] and John C. Doyle[2]*

Advanced technologies and biology have extremely different physical implementations, but they are far more alike in systems-level organization than is widely appreciated. Convergent evolution in both domains produces modular architectures that are composed of elaborate hierarchies of protocols and layers of feedback regulation, are driven by demand for robustness to uncertain environments, and use often imprecise components. This complexity may be largely hidden in idealized laboratory settings and in normal operation, becoming conspicuous only when contributing to rare cascading failures. These puzzling and paradoxical features are neither accidental nor artificial, but derive from a deep and necessary interplay between complexity and robustness, modularity, feedback, and fragility. This review describes insights from engineering theory and practice that can shed some light on biological complexity.

The theory and practice of complex engineering systems have progressed so radically that they often embody Arthur C. Clarke's dictum, "Any sufficiently advanced technology is indistinguishable from magic." Systems-level approaches in biology have a long history (1, 2) but are just now receiving renewed mainstream attention (3–13), whereas systems-level design has consistently been at the core of modern engineering, motivating its most sophisticated theories in controls, information, and computation. The hidden nature of complexity ("magic") and discipline fragmentation within engineering have been barriers to a dialog with biology. A key starting point in developing a conceptual and theoretical bridge to biology is robustness, the preservation of particular characteristics despite uncertainty in components or the environment (14).

Biologists and biophysicists new to studying complex networks often express surprise at a biological network's apparent robustness (15). They find that "perfect adaptation" and homeostatic regulation are robust properties of networks (16, 17), despite "exploratory mechanisms" that can seem gratuitously uncertain (18–20). Some even conclude that these mechanisms and their resulting features seem absent in engineering (20, 21). However, ironically, it is in the nature of their robustness and complexity that biology and advanced engineering are most alike (22). Good design in both cases (e.g., cells and bodies, cars and airplanes) means that users are largely unaware of hidden complexities, except through system failures. Furthermore, the robustness and fragility features of complex systems are both shared and necessary. Although the need for universal principles of complexity and corresponding mathematical tools is widely recognized (23), sharp differences arise as to what is fundamental about complexity and what mathematics is needed (24). This article sketches one possible view, using experience and theoretical insights from engineering complexity that are relevant to biology.

[1]Departments of Anesthesiology and Cell and Developmental Biology, University of Michigan Medical School, Ann Arbor, MI 48109, USA. [2]Control and Dynamical Systems, Electrical Engineering, and Bioengineering, California Institute of Technology, Pasadena, CA 91125, USA.

*To whom correspondence should be addressed. E-mail: doyle@cds.caltech.edu

We hope to dispel some common misconceptions and to renew a dialog between engineering theorists and their biologist and clinician colleagues.

Complexity, Optimality, and Convergent Evolution

The differences between biology and technology (and between organisms) are obvious, particularly at the molecular and device level. Nevertheless, convergent evolution, a well-established concept in both engineering and evolutionary biology, yields remarkable similarities at higher levels of organization. Recently, engineering systems have begun to have almost biological levels of complexity. For example, a Boeing 777 is fully "fly-by-wire" with 150,000 different subsystem modules, organized via elaborate protocols into complex control systems and networks, including roughly 1000 computers that can automate all vehicle functions. In terms of cost and complexity, the 777 is essentially a vast control system and computer network that just happens to fly. The consequence of good design is that its regulatory complexity is hidden from passengers (except when they use entertainment systems). The internal activity level is staggering, however (e.g., the data rate recorded on the internal state during final production testing is nearly equivalent to one human genome every minute). Commercial aircraft are not the only systems undergoing such explosions in complexity as a result of advanced controls and embedded networking; virtually all technologies are evolving similarly (25). We claim that this technological evolution of complexity is convergent with that of biology.

A striking example of convergent evolution is Fig. 1, comparing cruise speed to mass M over 12 orders of magnitude, from the 747 and 777 to fruit flies (26). The essential assumption in allometric scaling theory is that convergent evolution leads to nearly optimal systems with similar gross characteristics. It follows that simple arguments based on optimal design can explain functional relations between variables across many scales (27, 28). Here, a well-known elementary argument (29) shows good correspondence with the data and yields explanations for deviations. The popular allometric scaling theories (connecting, say, efficiency and geometry) are appealing: They are simple, accessible, suggestive evidence confirming convergent evolution and engineering optimality. Such theories are largely irrelevant to complexity directly, but an understanding of them leads to what is relevant. The scaling theory described by Fig. 1 does not distinguish between flight in the atmosphere and in a laboratory wind tunnel. In the latter context, a much simpler "mutant" 777 with nearly all of its 150,000-count "aeronome" knocked out would have roughly the same lift, mass, and cruise speed, and thus (from an allometric scaling viewpoint) would exhibit no deleterious laboratory "phenotype." Redundancy does not explain this finding (30). Rather, the mutant has lost control systems and robustness required for real flight outside the lab. Allometric scaling emphasizes the essential similarities between these 777 variants and a toy scale model (and a fruit fly), whereas our interest is their huge differences in complexity. Similarly, minimal cellular life requires a few hundred genes (31), yet even *Escherichia coli* have ~4000 genes, less than 300 of which have been classified as "essential" (32). The likely reason for this "excess" complexity is also the presence of complex regulatory networks for robustness. In technology as well as in organisms, such robustness tradeoffs drive the evolution of spiraling complexity.

As an example of spiraling complexity to battle fragility, consider our use of O_2 as a nutrient (electron acceptor), which obligates us to use complex feedback control mechanisms to ensure both sufficient O_2 and protection from O_2 toxicity. Distributed, multiscaled networks maintain precise internal, local O_2 concentrations throughout the body, both acutely and chronically. Dependency on such regulation makes its failure lethal, of course, but an additional fragility created by this exquisitely controlled environment is that it creates an attractive ecosystem for parasites, whose systems can thus be more streamlined. Host robustness to parasites then requires a separate complex immune control system.

In the developing immune system, T cells are educated to recognize self from nonself in the thymus. They are then selected to proceed on to the periphery (positive selection) or, if they are inaccurate sensors, they self-destruct. A fragility of this exceedingly complex immune system is autoimmune disease, an example being primary biliary cirrhosis (PBC) in which self-reactive lymphocytes slip through the immunity education program. Autoimmune injury to bile ducts causes toxic bile acids to accumulate. Injured hepatocytes fail to clear hormones, contributing to increased pressures in the liver circulation and even to rupture of connected venous systems. Pressure-induced distortion of blood vessels in the spleen traps platelets, and damaged hepatocytes undersynthesize blood-clotting proteins, both exaggerating blood loss after trauma. Unable to capitalize on the usual homeostatic feedback interactions with the liver, virtually every organ—including brain and kidney—can fail, all initiated by superficially minuscule autoimmune damage.

Medical interventions for PBC (i.e., drugs and transplantation) are further control systems adding to spiraling complexity, robustness, and fragility. Genetic variation in the P450 enzyme family leads to considerable interindividual variation in drug handling and side effects. Imbalance of the P450 network can lead to accumulation of toxins, including carcinogens. Polypharmacy—a common necessity—results in even more unpredictable interactions because drugs modulate P450 activity. Liver transplantation is now standard therapy for PBC. Immunosuppression must be sufficient to quash the immune mechanism that recognizes a foreign invader, but too much immunosuppression allows infection and tumors to go unchecked. Hence, the fragilities of transplantation are infection and tumors.

Modularity and Protocols

What emerges from these examples is that spiraling complexity, feedback regulation, robustness, fragility, and cascading failures are heavily intertwined, as is well known to biologists and engineers alike. Equally important and well known is the obvious role that modularity plays at every level, from base pairs and amino acids to genes and proteins, from organelles and membranes to pathways and networks, and finally to organs and organ axes (4–8)—and in every complex process, from development (11) to evolution (18). Although their meaning varies, modules generally are components, parts, or subsystems of a larger system that contain some or all of the following features: (i) identifiable interfaces (usually involving protocols) to other modules, (ii) can be modified and evolved somewhat independently, (iii) facilitate simplified or abstract modeling, (iv) maintain some identity when isolated or rearranged, yet (v) derive additional identity from the rest of the system.

The organization and design of advanced technologies suggest universal principles, relevant to biology, linking modularity with the robust yet fragile nature of complex systems. Truly universal principles should manifest
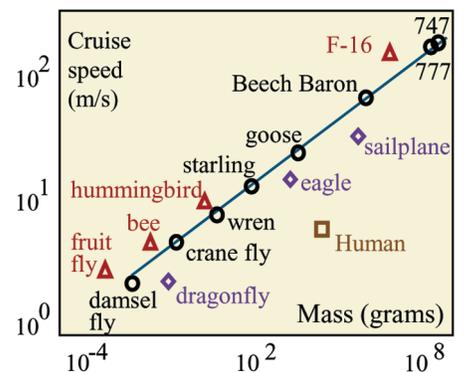
Fig. 1. Optimal cruise speed at sea level versus mass (log-log) for organisms and airplanes. Line is theoretical prediction (12) with $V = cM^\alpha$ and $\alpha = 1/6$ (29). Shorter wings for speed and maneuverability (triangles) yield higher cruise speeds than those optimized for soaring (diamonds). Most systems (circles) are compromises. Humans are not selected for powering flight and are far from optimal (square). Data and theory are from (26).

themselves in at least limited ways in scale-model (toy) systems, just as allometric scaling does. Consider the ubiquitous Lego toy system (*33, 34*). The signature feature of Lego is the patented snap connection for easy but stable assembly of components. The snap is the basic Lego protocol, and Lego bricks are its basic modules.

We claim that protocols are far more important to biologic complexity than are modules. They are complementary and intertwined but are important to distinguish. In everyday usage, protocols are rules designed to manage relationships and processes smoothly and effectively. If modules are ingredients, parts, components, subsystems, and players, then protocols describe the corresponding recipes, architectures, rules, interfaces, etiquettes, and codes of conduct (*35*). Protocols here are rules that prescribe allowed interfaces between modules, permitting system functions that could not be achieved by isolated modules. Protocols also facilitate the addition of new protocols and organization into collections of mutually supportive protocol suites. Like modules, they simplify modeling and abstraction, and as such may often be largely "in the eye of the beholder." A good protocol is one that supplies both robustness and evolvability.

Lego exhibits multilayer robustness, from components and toys to the product line. Lego bricks and toys are reusable and robust to trauma, and the snap is versatile, permitting endless varieties of toys from an array of components. This makes both a given Lego collection and the entire toy system evolvable to changes in what one chooses to build, to the addition of new Lego-compatible parts, and to novel toy designs. Evolution here is simply robustness to (possibly large) changes on long time scales. The low cost of modules and the popularity of the system confer other forms of robustness and evolvability; lost parts are easily replaced, and enthusiasts constantly design new modules and toys. The Lego protocol also creates fragilities at every level. Superficially minuscule damage to the snap at a key interface may cause an entire toy to fail, yet noninterfacing parts of bricks may be heavily damaged with minimal impact. The success of Lego means that any new snap, even a superior one, would not be easily adopted. Selection pressures thus preserve a protocol in two ways: Protocols facilitate evolution and are difficult to change.

It is instructive to compare the robustness properties (basic performance, ability to withstand trauma, versatility of allowed interconnections, reusability of modules, cost of parts and labor, and evolvability) of the standard Lego snap protocol (called the wild type, WT) with those of other hypothetical protocols (denoted *Smooth*, *Glue*, and *Mold*). Smooth bricks without snaps have unconstrained interconnections, but the results are

much less robust to trauma, severely limiting the range of toys. *Glue*, in addition to the WT snap, increases ability to withstand trauma but sharply decreases component reusability. Injection *Mold*ing entire toys goes even further. Thus, each "mutation" offers advantages, with both different robustness and fragility, but none uniformly improves on WT's overall robustness. WT is "fine-tuned" for robustness. We claim that this kind of optimality and robustness is most important to biological complexity.

As systems become more complex, protocols facilitate the layering of additional protocols, particularly involving feedback and signaling. Suppose we want to make a Lego structure incrementally more useful and versatile by "evolving" it to be (i) mobile, then (ii) motorized, then (iii) able to avoid collisions in a maze of obstacles. The first increment is easy to achieve, with Lego protocol–compatible axles and wheels. Motorizing toys involves a second increment in complexity, requiring protocols for motor and battery interconnection as well as a separate protocol for gears. All can be integrated into a motorized protocol suite to make modular subassemblies of batteries, motors, gears, axles, and wheels. These are available, inexpensive additions. The third increment increases cost and complexity by orders of magnitude, requiring layers of protocols and modules for sensing, actuation, and feedback controls plus subsidiary but essential ones for communications and computing (*34*). All are available, but it is here that we begin to see the true complexity of advanced technologies. Unfortunately, we also start to lose the easily described, intuitive story of the basic protocols. Minimal descriptions of advanced Lego features enabling sensing and feedback control literally fill books, but the protocols also facilitate the building of elaborate, robust toys, precisely because this complexity is largely hidden from users. This is consistent with the claim that biological complexity too is dominated not by minimal function, but by the protocols and regulatory feedback loops that provide robustness and evolvability.

This added complexity also creates new and often extreme fragilities. Removing a toy's control system might cause reversion to mere mobility, but a small change in an otherwise intact control system could cause wild, catastrophic behavior. For example, a small software bug might easily lead to collision seeking, a fragility absent in simpler toys. Similarly, large multicellular organisms are unaffected by the death of a single cell, but failure of one cell's control system can lead to fatal autoimmune diseases or cancer.

The snap protocol is concretely instantiated only in Lego modules, but it is also easy to identify the protocol itself as a useful and informative abstraction. The snap protocol is

more fundamental to Lego than are any individual modules. Similarly, we have no trouble distinguishing the many higher level protocols that organize sensing and feedback from the hardware modules themselves. In biology, the identification of protocols is easiest when shared by many different modules, as in Lego. Thus, abstractions such as gene regulation (*11*), covalent modification, membrane potentials, metabolic and signal transduction pathways, action potentials, and even transcription-translation, the cell cycle, and DNA replication could all be reasonably described as protocols (*36*), with their attendant modular implementations in various activators and repressors, kinases and phosphatases, ion channels, receptors, heterotrimeric guanine nucleotide binding proteins (G proteins), and so on. The cardiovascular system has protocols for gas and nutrient exchange and transport, implemented in heart, lung, vascular networks, and blood modules. The immune system involves elaborate protocols for complement and cell-mediated activation, implemented in modules such as T cells, natural killer cells, major histocompatibility complex molecules, and antibodies. Metazoan development has highly conserved protocols (*18*). Appropriate temporal and spatial expression during development (*11*) is regulated by enormous numbers of feedback strategies (*9*). These and many other protocols facilitate robust development and function in ways similar to Lego protocols, and they produce similar fragilities (*9*).

Thinking in terms of protocols, in addition to genes, organisms, and populations, as foci of natural selection, may be a useful abstraction for understanding the evolution of complexity (*37*). Good protocols allow new functions to be built from existing components and allow new components to be added or to evolve from existing ones, powerfully enhancing both engineering and evolutionary "tinkering." Protocols enable modularity and robustness but are in turn sources of fragility. Successful protocols become highly conserved because they both facilitate evolution and are difficult to change.

Lego has a perfectly complete "legome" of all parts, including full structure and function. A similar compendium is far from available for even simple organisms. Yet understanding a collision-avoiding, software-intensive, feedback-regulated Lego robot would require extensive reverse engineering of additional layers of protocols and modules beyond the legome. That the legome would not be sufficient is no surprise, but for reverse engineering such details may not be entirely necessary (see below). Imagine that such a Lego robot was a prototype for a single toy that dispensed entirely with the Lego modules in favor of custom implementation. Similar to *Mold*, this toy could easily have much more robustness to trauma, be faster, and navigate more complex obstacles, but at the

expense of limited part reuse. The modules and lower level protocols—most of the legome— would be completely different, yet we might claim that the essence of the toy, and what the prototype aimed to capture, remained. That essence involves the protocols that organized the sensors, actuators, and feedback control system that enables the obstacle avoidance and contributes almost the entire cost and complexity. These too are governed by protocols, but also by entirely new laws.

Elementary Feedback Concepts

Protocols are the most important aspect of modularity, and the most complex and critical protocols are for feedback control and the sensing, computing, communication, and actuation that implement it. Feedback control is both a powerful and dangerous strategy for creating robustness to external disturbances and internal component variations. Properly balanced, it delivers such a huge benefit that both engineers and evolution capitalize extensively on feedback to build and support complex systems (*4*, *9*). Detailed elaboration of the nature of regulatory feedback underlying complexity is beyond the scope of this article, but an elementary "toy" model illustrates the necessity of feedback to the function of complex systems as well as feedback's "conservation of fragility" law. This is arguably the most critical and rigorously established robustness tradeoff in complex systems.

In most technologies as well as in biochemistry, it is relatively easy to build either uncertain, high-gain components or precise, low-gain ones; but the precise, high-gain systems essential to both biology and technology are impossible or prohibitively expensive to make unless a feedback strategy like that in Fig. 2 is used. The simplest case is steady-state gain where, after some transient, r and d are held constant, and y too approaches a constant $y = Rr + Sd$ (*38*), where R and S are responses of y to s and d, respectively. Solving $y = d + ACy + Ar$ gives

$$y = ASr + Sd = \frac{1}{C}(S - 1)r + Sd$$

$$S \triangleq \frac{1}{1 - F} \quad F \triangleq AC \quad (1)$$

where F is the feedback gain. Ideally, perfect control would have $|S| = 0$, because that gives $y = -r/C$ ($R = -1/C$) completely independent of arbitrary variations in A and d. If $A \to \infty$ and $-1/C \gg 1$, then $F \to -\infty$, $|S| \to 0$, and $y \to -r/C$. Then R amplifies r and is perfectly robust to external disturbance d and to variations in A. Choosing C small and precise, with A sufficiently large and even sloppy, is one effective, efficient, and robust way to make y a high-gain function of r. $|S|$ measures the deviation from perfect control, and feedback can attenuate or greatly amplify the effects of un-

certainties. Defining fragility as $\log|S|$, note that $F < 0$ iff $|S| < 1$ iff $\log|S| < 0$ (*39*). $F > 0$ makes $\log|S| > 0$, amplifying d and uncertainty in A, and $F \to 1$ makes $\log|S| \to \infty$ (*40*). Unfortunately, this story is incomplete and even misleading without dynamics. The simplest possibility is for A and C to be first-order differential equations

$$C: \quad x' = -k_1 y - k_2 x \quad y = d + a$$
$$A: \quad a' = gu \qquad\qquad u = r + x \quad (2)$$

C is a low-pass filter with internal state x and parameters $k_1 > 0$ and $k_2 > 0$. A is a pure integrator with state a and gain $g > 0$ (*41*). This type of control is called "integral feedback." The parameters g, k_1, and k_2 might typically be functions of underlying physical quantities such as temperature, binding affinities, concentrations, etc., and thus might vary widely. The response $y(t)$ to steps in r and d are shown in Fig. 3 over two orders of magnitude in g and k_1. This simple protocol of integral feedback produces extremely robust external behavior even from wildly varying components (the blue solid versus red dashed lines in Fig. 3B). It is easily shown that this system is stable iff $gk_1 > 0$ and $k_2 > 0$, and converges to the steady state $y = (k_2/k_1)r$ independently of arbitrarily large variations in gain g and disturbance d (*42*). If $k_2 \gg k_1$, $y = (k_2/k_1)r$ is a high-gain amplifier as well (*43*). The individual values of g, k_1, and k_2 influence the rate of convergence to steady state, but only the ratio k_2/k_1 determines its value. Thus, robust high steady-state gain can be achieved with uncertain and small parameters with the right feedback protocol. Figure 3C shows that variations in both g and k_2 of orders of magnitude have modest impact, and only on early transient behavior.

The protocol here is the structure of the equations, including the integral feedback and the signs of the parameters. Modules are the implementations of the actuator and controller. As with Lego, the protocol must be "fine-tuned" (because rewiring components or flipping signs typically creates exponentially growing instabilities), but this allows the modules to
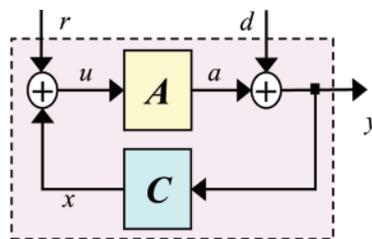

Fig. 2. Minimal feedback system with actuator A and controller/sensor C. The goal is for response y to amplify reference r, independent of external disturbance d and variations in A. The signals u and a are the input and output of the actuator A, and x is the output of C.

vary widely with minimal effect (*44*). Integral feedback is used ubiquitously in engineering (*45*) and is likely to be ubiquitous in biology as well, to achieve everything from homeostatic regulation to "perfect adaptation," and preliminary investigations confirm this impression (*46*–*48*). One reason is that integral feedback is both sufficient and necessary for perfect and robust steady-state tracking. Intuitively, necessity follows from the fact that in steady state, $a = y - d$ must perfectly cancel any constant (step in) d, whereas the input u to A cannot depend on this d, because y does not. Thus, A (or C) must contain an internal model of the dynamics of d, which for step changes is a pure integrator (*49*), which produces unbounded outputs to constant inputs. Thus, open-loop hypersensitivity is necessary for closed-loop robustness, and Fig. 3B is not an accident.

Fragility enters in the transient response. When g is increased, the response is faster but oscillatory (Fig. 3, A and B). Figure 3D plots fragility $\log|S(\omega)|$ versus ω where $Y(\omega)$ and $D(\omega)$ are Fourier transforms of y and d, and

$$\log|S(\omega)| = \log\left|\frac{Y(\omega)}{D(\omega)}\right| \quad (3)$$
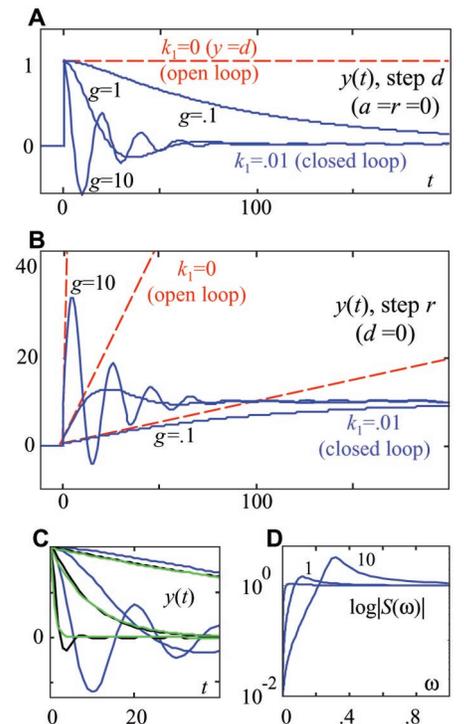

Fig. 3. Closed ($k_1 = 0.01$, blue) versus open ($k_1 = 0$, red) loop response $y(t)$ to step changes at $t = 0$ in (**A**) $d(t)$ ($r = 0$) and (**B**) $r(t)$ ($d = 0$) for $g = 0.1$, 1, and 10; $k_1 = 0.01$; and $k_2 = 10k_1$. Note the extreme divergence ($k_1 = 0$) versus convergence ($k_1 = 0.01$) as $t \to \infty$. (**C**) is a zoom of (A) with $k_1 = 0.01$, 0.1, and 1; $k_2 = 10k_1$ added for each value of g. (**D**) $\log|S(\omega)|$ versus ω for responses in (A). The peaks in $\log|S(\omega)|$ correspond to the oscillations in (A) and (B). Note the equal areas under the curves for $\log|S(\omega)|$.

For increasing g, low-frequency robustness ($\log|S(\omega)| < 0$) is improved, but at the expense of increased fragility ($\log|S(\omega)| > 0$) at higher frequencies (*50*). Indeed, it can be proven that for all g

$$\int_0^\infty \log|S(\omega)|\,d\omega = 0 \qquad (4)$$

so net fragility is, in this sense, a conserved quantity. Robustness ($\log|S(\omega)| < 0$) is paid for by an equal fragility ($\log|S(\omega)| > 0$), which amplifies d and uncertainty in A (*51*). This quite general result also holds for arbitrary parameters, control systems, and disturbances (*52*). Thus, there are always nonconstant (e.g., sinusoidal) $d(t)$ that would be amplified in $y(t)$. Such d could be perfectly rejected too, but only by adding internal models as complex as the external environment that generates d. Although such modeling is possible only for simple idealized laboratory environments, even approximate attempts can drive an extreme complexity spiral in real systems, and any controller is still subject to the constraint in Eq. 4. The key to good control design, then, is to ensure that this fragility is tolerable and occurs where uncertainties are relatively small.

Even these simple toy examples show the robust yet fragile features of complex regulatory networks. Their outward signatures are extremely constant regulated variables (yet occasional cryptic fluctuations) as well as extraordinary robustness to component variations (yet rare but catastrophic cascading failures). These apparently paradoxical combinations can easily be a source of confusion to experimentalists, clinicians, and theoreticians alike (*53*), but are intrinsic features of highly optimized feedback regulation. Because net robustness and fragility are constrained quantities, they must be manipulated and controlled with and within complex networks, even more so than energy and materials. Figure 3B shows how extreme open-loop versus closed-loop behavior can be, and thus how dangerous loss of control is to a system relying on it. The tradeoff in Eq. 4 shows that even when working perfectly, net fragility is constrained, and thus some transient amplification is unavoidable.

The necessity of integral feedback and the fragility constraint in Eq. 4 thus describe laws, not protocols—perhaps the two simplest such laws from control theory. Controllers that are more complex, with additional dynamics and multiple sensors and actuators, offer more refinement in performing robustness-fragility tradeoffs. Adding to regulatory complexity is also relatively easy in an evolutionary sense. Faster components allow for faster closed-loop responses. All are used in both biology and engineering, but all are still ultimately subject to Eq. 4. Control engineers must contend with this tradeoff, and its generalizations to more complex structures dominate control system design. Presumably, such tradeoffs dominate and constrain evolution and biology as well.

Implications for Biology and Engineering

The success of systems biology will certainly require modeling and simulation tools from engineering (*54, 55*), where experience shows that brute-force computational approaches are hopeless for complex systems involving protocols and feedback. Highly fragile features require highly sophisticated modeling, whereas robust features often have adequate models that are greatly simplified, requiring a "middle-out" approach (*10*). For example, if Fig. 2 is for a module in a larger system, the steady-state gain $y = (k_2/k_1)r$ depends only on k_2/k_1 and no other parameters, potentially simplifying experiments and modeling. If transient dynamics or component failure were of interest, more details would be needed, determined more by the rest of the system than by the internal components.

Many challenges of postgenomic biology are converging to the challenges facing engineers building complex "networks of networks," and engineering theory and practice are undergoing a revolution as radical as biology's. The simple ideas here only hint at the possibilities. For example, more complex control protocols than Fig. 2, used in both engineering and biology, can ameliorate (though not eliminate) the constraint in Eq. 4, but sophisticated theory is needed to elucidate the issues. Realistic models of biological networks will not be simple and will require multiple feedback signals, nonlinear component dynamics, numerous uncertain parameters, stochastic noise models (*56*), parasitic dynamics, and other uncertainty models. Scaling to deal with large networks will be a major challenge. Fortunately, researchers in robust control theory, dynamical systems, and related areas have been vigorously pursuing mathematics and software tools to address exactly these issues and apply them to complex engineering systems (*57, 58*). Biological applications are new, but progress so far is encouraging.

Experiments, modeling and simulation, and theory all have fragilities, but they are complementary, and through the right protocols they have the potential to create a robust "closed-loop" systems biology (*59*). Biologists' frustrating experience with theory has been primarily in an open-loop mode, where simple and attractive ideas can be wrong but receive enormous attention. Biology is the only science where feedback control and protocols play a dominant role, so it should not be surprising that there would be popular theories, coming from within science, that did not emphasize these issues. Biologists and engineers now have enough examples of complex systems that they can close the loop and eliminate specious theories (*60*). We should compare notes.

References and Notes

1. L. von Bertalanffy, *Modern Theories of Development: An Introduction to Theoretical Biology* (Oxford Univ. Press, New York, 1933).
2. M. A. Savageau, *Biochemical Systems Theory* (Addison-Wesley, Reading, MA, 1976).
3. _____, *Genetics* **149**, 1665 (1998).
4. C. V. Rao, A. P. Arkin, *Annu. Rev. Biomed. Eng.* **3**, 391 (2001).
5. M. Dickinson et al., *Science* **288**, 100 (2000).
6. D. Endy, R. Brent, *Nature* **409**, 391 (2001).
7. L. H. Hartwell, J. J. Hopfield, S. Leibler, A. W. Murray, *Nature* **402**, C47 (1999).
8. D. A. Lauffenburger, *Proc. Natl. Acad. Sci. U.S.A.* **97**, 5031 (2000).
9. M. Freeman, *Nature* **408**, 313 (2000).
10. D. Noble, *Science* **295**, 1678 (2002).
11. E. H. Davidson et al., *Science* **295**, 1669 (2002).
12. The supplements at *Science* Online (www.sciencemag.org/cgi/content/full/295/5560/1664/DC1) and at www.cds.caltech.edu/~doyle contain additional details, biology and engineering examples, proofs of math results, and references. Every topic considered here has a long history in biology (and engineering), renewed mainstream interest in the postgenomics era, and intense current research activity. Space limits force a focus on recent reviews that have better historical coverage than is possible here.
13. Alliance for Cellular Signaling (www.cellularsignaling.org).
14. Complex systems may be nonlinear, heterogeneous, large-scale, hierarchical, adaptive, etc., but these are subsidiary issues here.
15. Fragility here has the specific meaning of large and deleterious changes in particular system properties, attributable to possibly small but specific variations in the environment or in components. Robustness is often used more broadly, but here it means roughly the inverse of fragility. Their subtle interplay is the point of this article.
16. U. Alon, M. G. Surette, N. Barkai, S. Leibler, *Nature* **397**, 168 (1999).
17. G. von Dassow et al., *Nature* **406**, 188 (2000).
18. J. Gerhart, M. Kirschner, *Proc. Natl. Acad. Sci. U.S.A.* **95**, 8420 (1998).
19. Such mechanisms include chemotaxis, immune response, microtubule formation, angiogenesis, neurogenesis, and evolution itself. See (*18*).
20. M. Kirschner, J. Gerhart, T. Mitchison, *Cell* **100**, 79 (2000).
21. M. Chicurel, *Nature* **408**, 900 (2000).
22. This too has a long history; see (*1, 61*).
23. Special issue on Complex Systems, *Science* **284** (2 April 1999).
24. J. M. Carlson, J. Doyle, *Proc. Natl. Acad. Sci. U.S.A.* **99**, 2538 (2002).
25. HowStuffWorks, Inc. (www.howstuffworks.com).
26. H. Tennekes, *The Simple Science of Flight: From Insects to Jumbo Jets* (MIT Press, Cambridge, MA, 1997).
27. J. H. Brown, G. B. West, Eds., *Scaling in Biology* (Oxford Univ. Press, New York, 2000).
28. J. Whitfield, *Nature* **413**, 342 (2001).
29. Lift $L = \rho C_L A V^2$ must offset total weight W; A is wing area, ρ is air density, and C_L depends on geometry. Assuming fixed shapes and materials ($W^2 = C_5 A^3$) and eliminating A yields $V = CW^\alpha$ or $V = cM^\alpha$ where $\alpha = 1/6$ and c are constants. For more discussion of this greatly oversimplified treatment, see (*12, 26*).
30. Redundancy is a critical strategy for robustness to component failure, but it is already widely appreciated and adds modest complexity, so it is not discussed here.
31. Mycoplasma genitalium has 468 genes, and minimal sets for a free-living cell are estimated to be less than 300 genes (*62*).
32. "Essential" here refers to laboratory viability from single gene knockouts. See Profiling of E. coli Chromosome (PEC) (www.shigen.nig.ac.jp/ecoli/pec).
33. The LEGO Group (www.lego.com).
34. D. Baum, M. Gasperi, R. Hempel, L. Villa, *Extreme Mindstorms: An Advanced Guide to Lego Mindstorms* (Apress, Berkeley, CA, 2000).
35. "Laws" in society are protocols, unlike "laws" in science, which are necessary conditions that constrain (but are not subject to) design or selection (e.g., energy, mass,

fragility). A key research problem is distinguishing among laws, protocols, and historical accidents.

36. The central dogma can be thought of as a protocol, with DNA, RNA, protein, RNA polymerase, ribosome, etc., as modules.

37. A compelling case, but using very different terminology, is made in (*18*).

38. Steady state here means simply that all variables in Fig. 2 (*r*, *d*, *y*, *A*, *C*, etc.) approach constants, which can be solved for algebraically.

39. "iff" means "if and only if."

40. An important use of positive feedback is to deliberately destabilize equilibria and amplify small differences to create switches and to break symmetries and homogeneities. This can create patterns that are then maintained using negative feedback. Positive feedback is also critical to autocatalysis in growth and metabolism.

41. $da/dt = a' = gu$ means that a (the output of A) is a time integral of gu, where u is the input to A.

42. Stability is easily shown using standard methods of linear systems. Steady-state values can be found (in a stable system) by setting all time derivatives to 0, yielding $gk_1y = gk_2r$ or $y = (k_2/k_1)r$.

43. Mechanisms often exist that allow controller parameters (e.g., k_1 and k_2) to be much less uncertain than g and d. It is often even easier to make ratios such as k_2/k_1 largely invariant to variations in underlying physical quantities affecting the individual k_1 and k_2.

44. If precise gain is required, then the ratio k_2/k_1 must also be precise.

45. The national power grid has integral control at the >3000 power plants to regulate frequency and voltages of delivered power, oil refineries have >10,000 such control loops, and Internet congestion control involves a form implemented as part of TCP (transport control protocol). See (*12*) for more details, proofs, and examples.

46. T. M. Yi, Y. Huang, M. I. Simon, J. Doyle, *Proc. Natl. Acad. Sci. U.S.A.* **97**, 4649 (2000).

47. A. Levchenko, P. A. Iglesias, *Biophys. J.* **82**, 50 (2002).

48. H. El-Samad, J. P. Goff, M. Khammash, *J. Theor. Biol.* **214**, 17 (2002).

49. This argument can be made rigorous and is a standard elementary result in control theory. It is a special case of the internal model principle.

50. For sufficiently large g, the frequency domain peak and time domain transients become unacceptably large, although still stable.

51. One interpretation is that negative feedback is always balanced by an equal and opposite positive feedback. Strictly speaking, with dynamics this is not well defined, and $\log|S(\omega)|$ gives the correct generalization.

52. Relatively rare circumstances can involve an inequality ($\geq$). This is worse, but it means that Eq. 4 is an inequality constraint rather than a pure "conservation" law. See (*12, 63*).

53. The robust yet fragile nature of highly optimized complex regulatory networks can be mistakenly attributed to various kinds of bifurcations and "order-disorder" transitions (e.g., phase transitions, critical phenomena, "edge-of-chaos," pattern formation, etc.). See (*12, 24*).

54. The development of the Boeing 777 alone required a global software and computing infrastructure with roughly 10,000 workstations, terabytes of data, and a billion-dollar price tag.

55. Systems Biology Workbench (www.cds.caltech.edu/sbw).

56. D. T. Gillespie, *J. Chem. Phys.* **113**, 297 (2000).

57. K. Zhou, J. Doyle, K. Glover, *Robust and Optimal Control* (Prentice-Hall, Upper Saddle River, NJ, 1996).

58. E. Sontag, *Mathematical Control Theory: Deterministic Finite Dimensional Systems* (Springer, New York, 1998).

59. The idea of "closed-loop" systems biology is a purely metaphorical application of the message of this article to biology research itself. Biologists already iterate (loop) between experiment and at least informal modeling and "simulation." This process might be enhanced by "closing the loop" with theory and infrastructure "modules," which would also benefit.

60. One such example is Internet technology, which is rich in protocols and feedback and is beginning to have a rich theory. Even though it is poorly understood by nonexperts and has become a focus of many specious theories, details and enormous data sets are available, and it makes an attractive example to compare with biological networks. See (*12*).

61. N. Wiener, *Cybernetics* (Wiley, New York, 1948).

62. A. R. Mushegian, E. V. Koonin, *Proc. Natl. Acad. Sci. U.S.A.* **93**, 10268 (1996).

63. H. W. Bode, *Network Analysis and Feedback Amplifier Design* (Krieger, Melbourne, FL, 1945).

REVIEW

A Genomic Regulatory Network for Development

Eric H. Davidson,[1]* Jonathan P. Rast,[1] Paola Oliveri,[1] Andrew Ransick,[1] Cristina Calestani,[1] Chiou-Hwa Yuh,[1] Takuya Minokawa,[1] Gabriele Amore,[1] Veronica Hinman,[1] César Arenas-Mena,[1] Ochan Otim,[1] C. Titus Brown,[1] Carolina B. Livi,[1] Pei Yun Lee,[1] Roger Revilla,[1] Alistair G. Rust,[2]† Zheng jun Pan,[2]‡ Maria J. Schilstra,[2] Peter J. C. Clarke,[2] Maria I. Arnone,[3] Lee Rowen,[4] R. Andrew Cameron,[1] David R. McClay,[5] Leroy Hood,[4] Hamid Bolouri[2]

Development of the body plan is controlled by large networks of regulatory genes. A gene regulatory network that controls the specification of endoderm and mesoderm in the sea urchin embryo is summarized here. The network was derived from large-scale perturbation analyses, in combination with computational methodologies, genomic data, cis-regulatory analysis, and molecular embryology. The network contains over 40 genes at present, and each node can be directly verified at the DNA sequence level by cis-regulatory analysis. Its architecture reveals specific and general aspects of development, such as how given cells generate their ordained fates in the embryo and why the process moves inexorably forward in developmental time.

The mechanism causing cats to beget cats and fish to beget fish is hardwired in the genomic DNA, because the species specificity of the body plan is the cardinal heritable property. But despite all the examples of how individual genes affect the developmental process, there is yet no case where the lines of causality can be mapped from the genomic sequence to a major process of bilaterian development. One reason for this is that most of the developmental systems that have been intensively studied produce adult body parts, such as the third instar *Drosophila* wing disc, or the vertebrate hindbrain during rhombomere specification, or the heart anlagen of flies and mice (*1*). These systems present tough challenges because they go through successive stages of pattern formation in order to generate complex morphologies, and their development is initiated from states that are already complex. Furthermore, traditional molecular, genetic, and developmental biological approaches have focused on determining the functions of one or a few genes at a time, an approach that is not adequate for analysis of large regulatory control systems organized as networks. The heart of such networks consists of genes encoding transcription factors and the cis-regulatory elements that control the expression of those genes. Each of these cis-regulatory elements receives multiple inputs from other genes in the network; these inputs are the transcription factors for which the element contains the specific target site sequences. The functional linkages of which the network is composed are those between the outputs of regulatory genes and the sets of genomic target sites to which their products bind. Therefore, these linkages can be tested and verified by cis-regulatory analysis. This means identifying the control elements and their key target sites, and experimentally determining their functional significance. The view taken here is that "understanding" why a given developmental process

[1]Division of Biology, California Institute of Technology, Pasadena, CA 91125, USA. [2]Science and Technology Research Centre, University of Hertfordshire, AL10 9AB, UK. [3]Stazione Zoologica Anton Dohrn, 80121 Naples, Italy. [4]Institute for Systems Biology, Seattle, WA 98105, USA. [5]Department of Zoology, Duke University, Durham, NC 27708–0325, USA.

*To whom correspondence should be addressed. E-mail: davidson@caltech.edu
†Present address: European Bioinformatics Institute, Wellcome Trust Genome Campus, Hinxton, Cambridgeshire CB101 1SD, UK.
‡Present address: Altera European Technology Centre, Holmers Farm Way, High Wycombe, Buckinghamshire HP12 4XF, UK.