6.945 Adventures in Advanced Symbolic Programming

Spring 2009

Building Robust Systems
an essay

Gerald Jay Sussman
Massachusetts Institute of Technology

January 13, 2007

Abstract

It is hard to build robust systems: systems that have acceptable behavior over a larger class of situations than was anticipated by their designers. The most robust systems are evolvable: they can be easily adapted to new situations with only minor modification. How can we design systems that are flexible in this way?

Observations of biological systems tell us a great deal about how to make robust and evolvable systems. Techniques originally developed in support of symbolic Artificial Intelligence can be viewed as ways of enhancing robustness and evolvability in programs and other engineered systems. By contrast, common practice of computer science actively discourages the construction of robust systems.

Robustness

It is difficult to design a mechanism of general utility that does any particular job very well, so most engineered systems are designed to perform a specific job. General-purpose inventions, such as the screw fastener, when they appear, are of great significance. The digital computer is a breakthrough of this kind, because it is a universal machine that can simulate any other information-processing machine. We write software that configures our computers to effect this simulation for the specific jobs that we need done.

We have been designing our software to do particular jobs very well, as an extension of our past engineering practice. Each piece of software is designed to do a relatively narrow job. As the problem to be solved changes, the software must be changed. But small changes to the problem to be solved do not entail only small changes to the software. Software is designed too tightly for there to be much flexibility. As a consequence, systems do not evolve gracefully. They are brittle and must be replaced with entirely new designs as the problem domain evolves.[1] This is slow and expensive.

Our engineered systems do not have to be brittle. The Internet has adapted from a small system to one of global scale. Our cities evolve organically, to accommodate new business models, life styles, and means of transportation and communication. Indeed, from observation of biological systems we see that it is possible to build systems that can adapt to changes in the environment, both individually and as an evolutionary ensemble. Why is this not the way we design and build most software? There are historical reasons, but the main reason is that we don't know how to do this generally. At this moment it is an accident if a system turns out to be robust to changes in requirements.

Redundancy and degeneracy

Biological systems have evolved a great deal of robustness. One of the characteristics of biological systems is that they are redundant. Organs such as the liver and kidney are highly redundant: there is vastly more capacity

[1] Of course, there are some wonderful exceptions. For example, `EMACS` [26] is an extensible editor that has evolved gracefully to adapt to changes in the computing environment and to changes in its user's expectations. The computing world is just beginning to explore "engineered frameworks," for example, Microsoft's `.net` and Sun's `Java`. These are intended to be infrastructures to support evolvable systems.

than is necessary to do the job, so a person with a missing kidney suffers no obvious incapacity. Biological systems are also highly degenerate: there are usually many ways to satisfy a given requirement.[2] For example, if a finger is damaged, there are ways that the other fingers may be configured to pick up an object. We can obtain the necessary energy for life from a great variety of sources: we can metabolize carbohydrates, fats, and proteins, even though the mechanisms for digestion and for extraction of energy from each of these sources is quite distinct.

The genetic code is itself degenerate, in that the map from codons (triples of nucleotides) to amino acids is not one-to-one: there are 64 possible codons to specify only about 20 possible amino acids. [19] As a consequence, many point mutations do not change the protein specified by a coding region. This is one way variation can accumulate without obvious phenotypic consequences. If a gene is duplicated (not an uncommon occurrence), the copies may diverge silently, allowing the development of variants that may become valuable in the future, without interfering with current viability. In addition, the copies can be placed under different transcriptional controls.

Degeneracy is a product of evolution, and it certainly enables evolution. Probably degeneracy is itself selected for because only creatures that have significant amounts of degeneracy are sufficiently adaptable to allow survival as the environment changes. For example, suppose we have some creature (or engineered system) that is degenerate in that there are several very different interdependent mechanisms to achieve each essential function. If the environment changes (or the requirements change) so that one of the ways of achieving an essential function becomes untenable, the creature will continue to live and reproduce (the system will continue to satisfy its specifications). But the subsystem that has become inoperative is now open to mutation (or repair), without impinging on the viability (or current operation) of the system as a whole.

Engineered systems may incorporate some redundancy, in critical systems where the cost of failure is extreme. But they almost never intentionally incorporate degeneracy of the kind found in biological systems, except as a side effect of designs that are not optimal.[3]

[2] Although clear in extreme cases, the distinction biologists make between redundancy and degeneracy is fuzzy at the boundary. For more information see [8].

[3] Indeed, one often hears arguments against building flexibility into an engineered system. For example, in the philosophy of the computer language Python it is claimed: "There should be one—and preferably only one—obvious way to do it."[25] Science does

Exploratory Behavior

One of the most powerful mechanisms of robustness in biological systems is exploratory behavior.[4] The idea is that the desired outcome is produced by a generate-and-test mechanism. This organization allows the generator mechanism to be general and to work independently of the testing mechanism that accepts or rejects a particular generated result.

For example, an important component of the rigid skeleton that supports the shape of a cell is an array of microtubules. Each microtubule is made up of protein units that aggregate to form the microtubule. Microtubules are continually created and destroyed in a living cell; they are created growing out in all directions. However, only microtubules that encounter a stabilizer in the cell membrane are stable, thus supporting the shape determined by the positions of the stabilizers. So the mechanism for growing and maintaining a shape is relatively independent of the mechanism for specifying the shape. This mechanism partly determines the shapes of many types of cells in a complex organism, and it is almost universal in metazoans.

Exploratory behavior appears at all levels of detail in biological systems. The nervous system of a growing embryo produces a vastly larger number of neurons than will persist in the adult. Those neurons that find appropriate targets in other neurons, sensory organs, or muscles will survive and those that find no targets kill themselves. The hand is fashioned by production of a pad and deletion, by apoptosis, of the material between the fingers. [34] Our bones are continually being remodeled by osteoblasts (which build bone) and osteoclasts (which destroy bone). The shape and size of the bones is determined by constraints determined by their environment: the parts that they must be associated with, such as muscles, ligaments, tendons, and other bones.

Because the generator need not know about how the tester accepts or rejects its proposals, and the tester need not know how the generator makes its proposals, the two parts can be independently developed. This makes adaptation and evolution more efficient, because a mutation to one or the other of these two subsystems need not be accompanied by a complementary

not usually proceed this way: In classical mechanics, for example, one can construct equations of motion using Newtonian vectoral mechanics, or using a Lagrangian or Hamiltonian variational formulation.[30] In the cases where all three approaches are applicable they are equivalent, but each has its advantages in particular contexts.

[4] This thesis is nicely explored in the book of Kirschner and Gerhart.[17]

mutation to the other. However, this isolation is expensive because of the wasted effort of generation and rejection of failed proposals.

Indeed, generate and test is a metaphor for all of evolution. The mechanisms of biological variation are random mutations: modifications of the genetic instructions. Most mutations are neutral in that they do not directly affect fitness because of degeneracy in the systems. Natural selection is the test phase. It does not depend on the method of variation, and the method of variation does not anticipate the effect of selection.

There are even more striking phenomena: even in closely related creatures some components that end up almost identical in the adult are constructed by entirely different mechanisms in the embryo.[5] For distant relationships divergent mechanisms for constructing common structures may be attributed to "convergent evolution," but for close relatives it is more likely evidence for separation of levels of detail, in which the result is specified in a way that is somewhat independent of the way it is accomplished.

Compartments and localization

Every cell in our bodies is a descendant of a single zygote. All the cells have exactly the same genetic endowment (about 1GByte of ROM!). However there are skin cells, neurons, muscle cells, etc. The cells organize themselves to be discrete tissues, organs, and organ systems. This is possible because the way a cell differentiates and specializes depends on its environment. Almost all metazoans share homeobox genes, such as the Hox complex. Such genes produce an approximate coordinate system in the developing animal, separating the developing animal into distinct locales.[6] The locales provide context for a cell to differentiate. And information derived from contact with its neighbors produces more context that selects particular behaviors from

[5]The cornea of a chick and the cornea of a mouse are almost identical, but the morphogenesis of these two are not at all similar: the order of the morphogenetic events is not even the same. Bard [4] section 3.6.1 reports that having divergent methods of forming the same structures in different species is common. He quotes a number of examples. One spectacular case is that the frog *Gastrotheca riobambae* (recently discovered by delPino and Elinson [7]) develops ordinary frog morphology from an embryonic disk, whereas other frogs develop from an approximately spherical embryo.

[6]This is a very vague description of a complex process involving gradients of morphogens. I do not intend to get more precise here, as this is not a paper about biology, but rather about how biology informs engineering.

the possible behaviors that are available in its genetic program.[7]

This kind of organization has certain clear advantages. Flexibility is enhanced by the fact that the signaling among cells is permissive rather than instructive. That is, the behaviors of cells are not encoded in the signals; they are separately expressed in the genome. Combinations of signals just enable some behaviors and disable others. This weak linkage allows variation in the implementation of the behaviors that are enabled in various locales without modification of the mechanism that defines the locales. So systems organized in this way are evolvable in that they can accomodate adaptive variation in some locales without changing the behavior of subsystems in other locales.

Good engineering has a similar flavor, in that good designs are modular. Consider the design of a radio receiver. There are several grand "body plans" that have been discovered, such as direct conversion, TRF (tuned radio frequency), and superheterodyne. Each has a sequence of locales, defined by the engineering equivalent of a Hox complex, that patterns the system from the antenna to the output transducer. For example, a superheterodyne has the following locales:

$$\text{Antenna} : \text{RF} : \text{Converter} : \text{IF} : \text{Detector} : \text{AF} : \text{Transducer}$$

Each locale can be instantiated in many possible ways. The RF section may be just a filter, or it may be an elaborate filter and amplifier combination. Indeed, some sections may be recursively elaborated (as if the Hox complex were duplicated!) to obtain multiple-conversion receivers.

Of course, unlike biological mechanisms, in analog electronics the components are usually not universal in that each component can, in principle, act as any other component. But in principle there are universal electrical building blocks (programmable computer with analog interfaces for example!). For low-frequency applications one can build analog systems from such blocks. If each block had all of the code required to be any block in the system, but was specialized by interactions with its neighbors, and if there were extra unspecialized "stem cells" in the package, then we could imagine building self-reconfiguring and self-repairing analog systems.

[7]We have investigated some of the programming issues involved in this kind of development in our Amorphous Computing project.[2]

The structure of compartments in biological systems is supported by an elaborate infrastructure. One important component of this infrastructure is the ability to dynamically attach tags to materials being manipulated. For example, in eukaryotic cells, proteins are constructed with tags describing their destinations. [19] For example, a transmembrane protein, which may be part of an ion-transport pore, is directed to the plasma membrane, whereas some other protein might be directed to the golgi apparatus. There are mechanisms in the cell (themselves made up of assemblies of proteins) to recognize these tags and effect the transport of the parts to their destinations. Tagging is also used to clean up and dispose of various wastes, such as protein molecules that did not fold correctly or that are no longer needed. Such proteins are tagged (ubiquinated) and carried to a proteasome for degradation.

This structure of compartments is also supported at higher levels of organization. There are tissues that are specialized to become boundaries of compartments, and tubes that interconnect them. Organs are bounded by such tissues and interconnected by such tubes, and the entire structure is packaged to fit into coelems, which are cavities lined with specialized tissues in higher organisms.

Defense, repair, and regeneration

Biological systems are always under attack from predators, parasites, and invaders. They have developed elaborate systems of defense, ranging from restriction enzymes[8] in bacteria to the immune systems of mammals. Advanced systems depend on continuous monitoring of the external and internal environment, and mechanisms for distinguishing self from other.

In a complex organism derived from a single zygote every cell is, in principle, able to perform the functions of every other cell. Thus there is redundancy in numbers. But even more important is the fact that this provides a mechanism for repair and regeneration. A complex organism is a dynamically reconfigurable structure made out of potentially universal interchange-

[8] A restriction enzyme cuts DNA molecules at particular sites that are not part of the genome of the bacterium, thus providing some defense against viruses that may contain such sites in their genome. One could imagine an analogous computational engine that stops any computation containing a sequence of instructions that does not occur in the code of the operating system. Of course, a deliberately constructed virus (biological or computational) may be designed to elude any such simple restriction-enzyme structure.

able and reproducible parts: if a part is damaged, nearby cells can retarget to fill in the gap and take on the function of the damaged part.

The computer software industry has only recently begun to understand the threats from predators, parasites, and invaders. The early software systems were built to work in friendly, safe environments. But with the globalization of the network and the development of economic strategies that depend on attacking and coopting vulnerable systems, the environment has changed to a substantially hostile one. Current defenses, such as antivirus and antispam software, are barely effective in this environment (although there are significant attempts to develop biologically-inspired computer "immune systems").

One serious problem is monoculture. Almost everyone uses the same computer systems, greatly increasing the vulnerability. In biological systems there are giant neutral spaces, allowing great variation with negligible change in function: humans can have one of several blood types; there are humans of different sizes, shapes, colors, etc. But they are all human. They all have similar capabilities and can all live in a great variety of environments. They communicate with language! However, not all humans have the same vulnerabilities: people heterozygous for the sickle-cell trait have some resistance to malaria.

In our engineered systems we have not, in general, taken advantage of this kind of diversity. We have not yet made use of the diversity that is available in alternate designs, or even used the variation that is available in silent mutations. Part of the reason is that there is economy in monoculture. But this economy is short-sighted and illusory, because of the extreme vulnerability of monoculture to deliberate and evolved attack.

Biological systems have substantial abilities to repair damage, and, in some cases, to regenerate lost parts. This ability requires extensive and continuous self-monitoring, to notice the occurrence of damage and initiate a repair process. It requires the ability to mobilize resources for repair and it requires the information about how the repair is to be effected.

Systems that build structure using exploratory behavior can easily be coopted to suppport repair and regeneration. However, it is still necessary to control the exploratory proliferation to achieve the desired end state. This appears to be arranged with homeostatic constraint mechanisms. For example, a wound may require the production of new tissue to replace the lost material. The new tissue needs to be supplied with oxygen and nutrients, and it needs wastes removed. Thus it must be provided with new capillar-

ies that correctly interconnect with the circulatory system. Cells that do not get enough oxygen produce hormones that stimulate the proliferation of blood vessels in their direction. Thus, the mechanisms that build the circulatory system need not know the geometry of the target tissues. Their critical mission is achieved by exploration and local constraint satisfaction. Such mechanisms support both morphogenesis in the embryo and healing in the adult.

There are very few engineered systems that have substantial ability for self-repair and regeneration. High-quality operating systems have "file-system salvagers" that check the integrity of the file system and use redundancy in the file system to repair broken structures and to regenerate some lost parts. But this is an exceptional case. How can we make this kind of self-monitoring and self-repair the rule rather than the exception?

In both cases, defense and repair, a key component is awareness—the ability to monitor the environment for imminent threats and one's self for damage.

Composition

Large systems are composed of many smaller components, each of which contributes to the function of the whole either by directly providing a part of that function or by cooperating with other components by being interconnected in some pattern specified by the system architect to establish a required function. A central problem in system engineering is the establishment of interfaces that allow the interconnection of components so that the functions of those components can be combined to build compound functions.

For relatively simple systems the system architect may make formal specifications for the various interfaces that must be satisfied by the implementers of the components to be interconnected. Indeed, the amazing success of electronics is based on the fact that it is feasible to make such specifications and to meet them. High-frequency analog equipment is interconnected with coaxial cable with standardized impedance characteristics, and with standardized families of connectors. [3] Both the function of a component and its interface behavior can usually be specified with only a few parameters. [14] In digital systems things are even more clear. There are static specifications of the meanings of signals (the digital abstraction). There are dynamic specifications of the timing of signals. [32] And there are mechanical specifications

of the form-factors of components.[9]

Unfortunately, this kind of *a priori* specification becomes progressively more difficult as the complexity of the system increases.[10] The specifications of computer software components are often enormously complicated. They are difficult to construct and it is even more difficult to guarantee compliance with such a specification. Many of the fragilities associated with software are due to this complexity.

By contrast, biology constructs systems of enormous complexity without very large specifications. The human genome is about 1 GByte. This is vastly smaller than the specifications of a major computer operating system. How could this possibly work? We know that the various components of the brain are hooked together with enormous bundles of neurons, and there is nowhere near enough information in the genome to specify that interconnect in any detail. What is likely is that the various parts of the brain learn to communicate with each other, based on the fact that they share important experiences. So the interfaces must be self-configuring, based on some rules of consistency, information from the environment, and extensive exploratory behavior. This is pretty expensive in boot-up time (it takes some years to configure a working human) but it provides a kind of robustness that is not found in our engineered entities to date.

The Cost

> "To the optimist, the glass is half full. To the pessimist, the glass is half empty. To the engineer, the glass is twice as big as it needs to be."

[9] *The TTL Data Book for Design Engineers* [31] is a classic example of a successful set of specifications for digital-system components. It specifies several internally consistent "families" of small-scale and medium-scale integrated-circuit components. The families differ in such characteristics as speed and power dissipation, but not in function. The specification describes the static and dynamic characteristics of each family, the functions available in each family, and the physical packaging for the components. The families are cross consistent as well as internally consistent in that each function is available in each family, with the same packaging and a consistent nomenclature for description. Thus a designer may design a compound function and later choose the family for implementation. Every good engineer (and biologist!) should be familiar with the lessons in the TTL Data Book.

[10] We could specify that a chess-playing program plays a legal game—that it doesn't cheat, but how would one begin to specify that it plays a good game of chess?

author unknown

Unfortunately, generality and evolvability require redundancy, degeneracy, and exploratory behavior. These are expensive, when looked at in isolation. A mechanism that works over a wide range of inputs must do more to get the same result as a mechanism specialized to a particular input. A redundant mechanism has more parts than an equivalent non-redundant mechanism. A degenerate mechanism appears even more extravagant. Yet these are ingredients in evolvable systems. To make truly robust systems we must be willing to pay for what appears to be a rather elaborate infrastructure. The value, in enhanced adaptability, may be even more extreme. Indeed, the cost of our brittle infrastructure probably greatly exceeds the cost of a robust design, both in the cost of disasters and in the lost opportunity costs due to the time of redesign and rebuilding.

The problem with correctness

But there may be an even bigger cost to building systems in a way that gives them a range of applicablity greater than the set of situations that we have considered at design time. Because we intend to be willing to apply our systems in contexts for which they were not designed, we cannot be sure that they work correctly!

We are taught that the "correctness" of software is paramount, and that correctness is to be achieved by establishing formal specification of components and systems of components and by providing proofs that the specifications of a combination of components are met by the specifications of the components and the pattern by which they are combined. I assert that this discipline enhances the brittleness of systems. In fact, to make truly robust systems we must discard such a tight discipline.

The problem with requiring proofs is that it is usually harder to prove general properties of general mechanisms than it is to prove special properties of special mechanisms used in constrained circumstances. This encourages us to make our parts and combinations as special as possible so we can simplify our proofs.

I am not arguing against proofs. They are wonderful when available. Indeed, they are essential for critical system components, such as garbage collectors (or ribosomes!). However, even for safety-critical systems, such as autopilots, the restriction of applicability to situations for which the system is

provably correct as specified may actually contribute to unnecessary failure. Indeed, we want an autopilot to make a good-faith attempt to safely fly an airplane that is damaged in a way not anticipated by the designer!

I am arguing against the discipline of *requiring* proofs: The requirement that everything must be proved to be applicable in a situation before it is allowed to be used in that situation excessively inhibits the use of techniques that could enhance the robustness of designs. This is especially true of techniques that allow a method to be used, on a tight leash, outside of its proven domain, and techniques that provide for future expansion without putting limits on the ways things can be extended.

Unfortunately, many of the techniques I advocate make the problem of proof much more difficult, if not practically impossible. On the other hand, sometimes the best way to attack a problem is to generalize it until the proof becomes simple.

Infrastructure to Support Generalizability

We want to build systems that can be easily generalized beyond their initial use. Let's consider techniques that can be applied in software design. I am not advocating a grand scheme or language, such as Planner, but rather infrastructure for integrating each of these techniques when appropriate.

Generality of parts

The most robust systems are built out of families of parts, where each part is of very general applicability. Such parts have acceptable behavior over a much wider class of conditions than is needed for any particular application. If, for example, we have parts that produce outputs for given inputs, we need the range of acceptable inputs for which the outputs are sensible to be very broad. To use a topological metaphor, the class of acceptable inputs for any component used in a solution to a current problem should be an "open set" surrounding the inputs it will encounter in actual use in the current problem.

Furthermore, the range of outputs of the parts over this wide range of inputs should be quite small and well defined: much smaller than the range of acceptable inputs. This is analogous to the static discipline in the digital abstraction that we teach to students in introductory computer systems subjects.[32] The power of the digital abstraction is that the outputs are al-

ways better than the acceptable inputs, so it suppresses noise. Using more general parts builds a degree of flexibility into the entire structure of our systems. Small perturbations of the requirements can be adjusted to without disaster, because every component is built to accept perturbed (noisy) inputs.

There are a variety of techniques to help us make families of components that are more general than we anticipate needing for the applications under consideration at the time of the design of the components. These techniques are not new. They are commonly used, often unconsciously, to help us conquer some particular problem. However, we have no unified understanding or common infrastructure to support their use. Furthermore, we have developed a culture that considers many of these techniques dangerous or dirty. It is my intention to expose these techniques in a unified context so we can learn to exploit them to make more robust systems.

Extensible generic operations

One good idea is to build a system on a substrate of extensible generic operators. Modern dynamically typed programming languages usually have built-in arithmetic that is generic over a variety of types of numerical quantities, such as integers, floats, rationals, and complex numbers.[28, 15, 24] This is already an advantage, but it surely complicates reasoning and proofs and it makes the implementation much more complicated and somewhat less efficient than simpler systems. However, I am considering an even more general scheme, where it is possible to define what is meant by addition, multiplication, etc., for new datatypes unimagined by the language designer. Thus, for example, if the arithmetic operators of a system are extensible generics, a user may extend them to allow arithmetic to be extended to quaternions, vectors, matrices, integers modulo a prime, functions, tensors, differential forms, ... This is not just making new capabilities possible; it also extends old programs, so a program that was written to manipulate simple numerical quantities may become useful for manipulating scalar-valued functions.[11]

[11] A mechanism of this sort is implicit in most "object-oriented languages," but it is usually buried in the details of ontological mechanisms such as inheritance. The essential idea of extensible generics appears in SICP [1] and is usefully provided in tinyCLOS [18] and SOS [12]. A system of extensible generics, based on predicate dispatching, is used to implement the mathematical representation system in SICM [30]. A nice exposition of predicate dispatching is given by Ernst [9].

However, there is a risk. A program that depends on the commutativity of numerical multiplication will certainly not work correctly for matrices. (Of course, a program that depends on the exactness of operations on integers will not work correctly for inexact floating-point numbers either.) This is exactly the risk that comes with evolution—some mutations will be fatal! But that risk must be balanced against the cost of not trying to use the program, in a pinch.

On the other hand, some mutations will be extremely valuable. For example, it is possible to extend arithmetic to symbolic quantities. The simplest way to do this is to make a generic extension to all of the operators to take symbolic quantities as arguments and return a data structure representing the indicated operation on the arguments. With the addition of a simplifier of algebraic expressions we suddenly have a symbolic manipulator. This is very useful in debugging purely numerical calculations, because if they are given symbolic arguments we can examine the resulting symbolic expressions to make sure that the program is calculating what we intend it to. It is also the basis of a partial evaluator for optimization of numerical programs. And functional differentiation can be viewed as a generic extension of arithmetic to a hyperreal datatype.[12]

Extensible generic operations are not for the faint of heart. The ability to extend operators *after the fact* gives both extreme flexibility and whole new classes of bugs! It is probably impossible to prove very much about a program when the primitive operations can be extended, except that it will work when restricted to the types it was defined for. This is an easy but dangerous path for generalization.

Extensible generic operations, and the interoperation of interpreted and compiled code, imply that data must be tagged with the information required to decide which procedures are to be used for implementing the indicated operations. But once we have the ability to tag data there are other uses tags can be put to. For example, we may tag data with its provenance, or how it was derived, or the assumptions it was based on. Such audit trails may be essential for access control, for tracing the use of sensitive data, or for debugging complex systems.[33] Thus we can get power by being able to attach arbitrary tags to any data item, besides the tags used for determining generics.

[12]The `scmutils` system we use to teach classical mechanics [30] implements differentiation in exactly this way.

Generate and test

We normally think of generate and test, and its extreme use in search, as an AI technique. However, it can be viewed as a way of making systems that are modular and independently evolvable, as in the exploratory behavior of biological systems. Consider a very simple example: suppose we have to solve a quadratic equation. There are two roots to a quadratic. We could return both, and assume that the user of the solution knows how to deal with that, or we could return one and hope for the best. (The canonical SQRT routine returns the positive square root, even though there are two square roots!) The disadvantage of returning both solutions is that the receiver of that result must know to try his computation with both and either reject one, for good reason, or return both results of his computation, which may itself have made some choices. The disadvantage of returning only one solution is that it may not be the right one for the receiver's purpose.

A better way to handle this is to build a backtracking mechanism into the infrastructure.[10, 13, 21, 1] The square-root procedure should return one of the roots, with the option to change its mind and return the other one if the first choice is determined to be inappropriate by the receiver. It is, and should be, the receiver's responsibility to determine if the ingredients to its computation are appropriate and acceptable. This may itself require a complex computation, involving choices whose consequences may not be apparent without further computation, so the process is recursive. Of course, this gets us into potentially deadly exponential searches through all possible assignments to all the choices that have been made in the program. As usual, modular flexibility can be dangerous.

But if the choice mechanism attaches a tag describing its state to the data it selects, and if the primitive operations that combine data combine these tags correctly, one can always tell which choices contributed to any particular piece of data. With such a system, search can be optimized so that only relevant choices must be considered in any particular backtrack. This is the essence of dependency-directed backtracking. [27, 11, 20, 29] If such a system is built into the infrastructure then exploratory behavior can be as efficient as any explicit manipulation of sets of choices, without a program having to know which contributors of data to its inputs are actually sets of possibilities. It does, however, incur the overhead of a program testing for consistency of its results and rejecting them if necessary. Of course, this is important in any system intended to be reliable as well as robust.

Constraints generalize procedures

Consider an explicit integrator for a system of ordinary differential equations, such as the Bulirsch-Stoer algorithm.[6, 23] The description of the ODEs for such an integrator is a system-derivative procedure that takes a state of the system and gives back the derivative of the state. For example, the system derivative for a driven harmonic oscillator takes a structure with components the time, the position, and the velocity and returns a vector of 1 (dt/dt), the velocity, and the acceleration. The system derivative has three parameters: the damping constant, the square of the undamped oscillatory frequency, and the drive function. The natural frequencies are determined by a quadratic in the first two parameters. The sum of the natural frequencies is the damping constant and the product of the natural frequencies is the square of the undamped oscillatory frequency. We can also define a Q for such a system. In any particular physical system, such as a series-resonant circuit, there are relationships between these parameters and the inductance, the capacitance, and the resistance of the circuit. Indeed, one may specify the system derivative in many ways, such as the oscillatory frequency, the capacitance, and the Q. There is no reason why this should be any harder than specifying the inductance, the resistance, and the capacitance.

If we have a set of quantities and relations among them we can build a constraint network that will automatically derive some of the quantities given the values of others. By including the parameters of the system derivative in a constraint network we greatly increase the generality of its application without any loss of efficiency for numerical integration. An added benefit is that we can use the constraint-propagation process to give us multiple alternative views of the mechanism being simulated: we can attach a spring, mass, and dashpot constraint network to our series RLC constraint network and think about the inertia of the current in our inductor. The infrastructure needed to support such a constraint-directed invocation mechanism is inexpensive, and the truth-maintenance system needed to track the dependencies is the same mechanism needed to implement the dependency-directed backtracking described above.

But constraints give us more than a support for generality. Constraints that dynamically test the integrity of data structures and hardware can be used to notice and signal damage. Such mechanisms may be able to encapsulate damage so that it does not spread, and trigger defense mechanisms to fight the damaging agent. Also, if we make systems that build themselves

using generate-and-test mechanisms controlled by constraints that enforce requirements on the structure of the result we can build systems that can repair some forms of damage automatically.

Degeneracy in engineering

In the design of any significant system there are many implementation plans proposed for every component at every level of detail. However, in the system that is finally delivered this diversity of plans is lost and usually only one unified plan is adopted and implemented. As in an ecological system, the loss of diversity in the traditional engineering process has serious consequences.

We rarely build degeneracy into programs, partly because it is expensive and partly because we have no formal mechanisms for mediating its use. However, there is a mechanism from the AI problem-solving world for degenerate designs: goal-directed invocation. The idea is that instead of specifying "how" we want a goal accomplished, by naming a procedure to accomplish it, we specify "what" we want to accomplish, and we link procedures that can accomplish that goal with the goal. This linkage is often done with pattern matching, but that is accidental rather than essential.

If there is more than one way to accomplish the goal, then the choice of an appropriate procedure is a choice point that can be registered for backtracking. Of course, besides using a backtrack search for choosing a particular way to accomplish a goal there are other ways that the goal can invoke degenerate methods. For example, we may want to run several possible ways to solve a problem in parallel, choosing the one that terminates first.

Suppose we have several independently implemented procedures all designed to solve the same (imprecisely specified) general class of problems. Assume for the moment that each design is reasonably competent and actually works correctly for most of the problems that might be encountered in actual operation. We know that we can make a more robust system by combining the given procedures into a larger system that independently invokes each of the given procedures and compares their results, choosing the best answer on every problem. If the combination has independent ways of determining which answers are acceptable we are in very good shape. But even if we are reduced to voting, we get a system that can reliably cover a larger space of solutions. Furthermore, if such a system can automatically log all cases where one of the designs fails, the operational feedback can be used to improve the performance of the procedure that failed.

This degenerate design strategy can be used at every level of detail. Every component of each subsystem can itself be so redundantly designed and the implementation can be structured to use the redundant designs. If the component pools are themselves shared among the subsystems, we get a controlled redundancy that is quite powerful. However, we can do even better. We can provide a mechanism for consistency checking of the intermediate results of the independently designed subsystems, even when no particular value in one subsystem exactly corresponds to a particular value in another subsystem.

For a simple example, suppose we have two subsystems that are intended to deliver the same result, but computed in completely different ways. Suppose that the designers agree that at some stage in one of the designs, the product of two of the variables in that design must be the same as the sum of two of the variables in the other design.[13] There is no reason why this predicate should not be computed as soon as all of the four values it depends upon become available, thus providing consistency checking at run time and powerful debugging information to the designers. This can be arranged using a locally embedded constraint network.

Again, if we make systems that build themselves using generate-and-test mechanisms controlled by constraints that enforce requirements on the structure of the result, we will get significant natural degeneracy, because there will in general be multiple proposals that are accepted by the constraints. Also, because of the environmental differences among the instances of the systems to be built we will automatically get variation from system instance to system instance. This neutral space variation will give substantial resistance to invasion.

Infrastructure to Support Robustness and Evolvability

Combinators

If the systems we build are made up from members of a family of mix-and-match components that combine to make new members of the family (by

[13]This is actually a real case: in variational mechanics the sum of a Lagrangian for a system and the Hamiltonian related to it by a Legendre transformation is the inner product of the generalized momentum 1-form and the generalized velocity vector.[30]

obeying a predetermined standard protocol of interconnect), bigger perturbations of the requirements are more easily addressed by rearrangement of high-level components.

But how do we arrange to build our systems by combining elements of a family of mix-and-match components? One method is to identify a set of primitive components and a set of combinators that combine components so as to make compound components with the same interface as the primitive components. Such sets of combinators are sometimes explicit, but more often implicit, in mathematical notation.

The use of functional notation is just such a discipline. A function has a domain, from which its arguments are selected, and a range of its possible values. There are combinators that produce new functions as combinations of others. For example, the composition of functions f and g is a new function that takes arguments in the domain of g and produces values in the range of f. If two functions have the same domain and range, and if arithmetic is defined on their common range, then we can define the sum (or product) of the functions as the function that when given an argument in their common domain, is the sum (or product) of the values of the two functions at that argument. Languages that allow first-class procedures provide a mechanism to support this means of combination, but what really matters is a good family of pieces.

There are entire families of combinators that we can use in programming that we don't normally think of. Tensors are an extension of linear algebra to linear operators with multiple arguments. But the idea is more general than that: the "tensor combination" of two procedures is just a new procedure that takes a data structure combining arguments for the two procedures. It distributes those arguments to the two procedures, producing a data structure that combines the values of the two procedures. The need to unbundle a data structure, operate on the parts separately, and rebundle the results is ubiquitous in programming. There are many such common patterns. It is to our advantage to expose and abstract these into a common library of combinators.

We may use constraints when we model physical systems. Physical systems have conservation laws that are expressed in terms of dual variables, such as torque and angular velocity, or voltage and current. Primitive constraints and combinators for such systems are a bit more complex, but some have been worked out in detail. For example, the wonderful MARTHA system of Penfield gives a complete set of combinators for electrical circuits

represented in terms of parts with two-terminal ports. [22]

Continuations

There are now computer languages that provide access to first-class continuations. This may seem to be a very esoteric construct, when introduced in isolation, but it enables a variety of control structures that can be employed to substantially improve the robustness of systems.

A continuation is a captured control state of a computation.[14] If a continuation is invoked, the computation continues at the place represented by the continuation. A continuation may represent the act of returning a value of a subexpression to the evaluation of the enclosing expression. The continuation is then a procedure that when invoked returns its argument to the evaluation of the enclosing expression as the value of the subexpression. A continuation is a first-class object that can be passed as an argument, returned as a value, and incorporated into a data structure. It can be invoked multiple times, allowing a computation to be resumed at a particular point with different values returned by the continuation.

Continuations give the programmer explicit control over time. A computation can be captured and suspended at one moment and restored and continued at any future time. This immediately provides coroutines (cooperative multitasking), and with the addition of a timer interrupt mechanism we get timesharing (preemptive multitasking).

Backtracking and concurrency

Continuations are a natural mechanism to support backtracking. A choice can be made, and if that choice turns out to be inappropriate, an alternative choice can be made and its consequences worked out. (Wouldn't we like real life to have this feature!) So, in our square-root example, the square-root program should return the AMB of both square roots, where AMB is the operator that chooses and returns one of them, with the option to provide the other if the first is rejected. The receiver can then just proceed to use the given solution, but if at some point the receiver finds that its computation

[14]This control state is not to be confused with the full state of a system. The full state is all the information required, along with the program, to determine the future of a computation. It includes all of the current values of mutable variables and data. The continuation does not capture the current values of mutable variables and data.

does not meet some constraint it can FAIL, causing the AMB operator to revise its choice and return with the new choice through its continuation. In essence, the continuation allows the generator of choices to coroutine with the receiver/tester of the choices.

If there are multiple possible ways to solve a subproblem, and only some of them are appropriate for solving the larger problem, sequentially trying them as in generate-and-test is only one way to proceed. For example, if some of the choices lead to very long (perhaps infinite) computations in the tester while others may succeed or fail quickly, it is appropriate to allocate each choice to a thread that may run concurrently. This requires a way for threads to communicate and perhaps for a successful thread to kill its siblings. All of this can be arranged with continuations, with the thread-to-thread communications organized around transactions.

Arbitrary association

The ability to annotate any piece of data with other data is a crucial mechanism in building robust systems. The attachment of metadata is a generalization of the tagging used to support extensible generic operations. Additional tags, labeling the data with the choices leading to its derivation, may be used to support dependency-directed backtracking. Sometimes it is appropriate to attach a detailed audit history, describing the derivation of a data item, to allow some later process to use the derivation for some purpose or to evaluate the validity of the derivation for debugging. For many missions, such as legal arguments, it is necessary to know the provenance of data: where it was collected, how it was collected, who collected it, how the collection was authorized, etc. The detailed derivation of a piece of evidence, giving the provenance of each contribution, may be essential to determining if it is admissable in a trial.

Associations of data items can be implemented by many mechanisms, such as hash tables. But the implementation may be subtle. For example, the cost of a product will in general depend on different assumptions than the shipping weight of the product, which may have the same numerical value. If the computational system does not have a different token for each of these two equal numbers, the system does not have a way of hanging distinct tags on them.

Dynamically configured interfaces

How can entities talk when they don't share a common language? A computational experiment by Simon Kirby has given us an inkling of how language may have evolved. In particular, Kirby [16] showed, in a very simplified situation, that if we have a community of agents that share a few semantic structures (perhaps by having common perceptual experiences) and that try to make and use rules to parse each other's utterances about experiences they have in common, then the community eventually converges so that the members share compatible rules. While Kirby's experiment is very primitive, it does give us an idea about how to make a general mechanism to get disparate modules to cooperate.

Jacob Beal [5] extended and generalized the work of Kirby. He built and demonstrated a system that allowed computational agents to learn to communicate with each other through a sparse but uncontrolled communication medium. The medium has many redundant channels, but the agents do not have an ordering on the channels, or even an ability to name them. Nevertheless, employing a coding scheme reminiscent of Calvin Mooers's Zatocoding (an early kind of hash coding), where descriptors of the information to be retrieved are represented in the distribution of notches on the edge of a card, Mr. Beal exchanges the sparseness and redundancy of the medium for reliable and reconfigurable communications of arbitrary complexity. Beal's scheme allows multiple messages to be communicated at once, by superposition, because the probability of collision is small. Beal has shown us new insights into this problem, and the results may be widely applicable to engineering problems.

Conclusion

Serious engineering is only a few thousand years old. Our attempts at deliberately producing very complex robust systems are immature at best. We have yet to glean the lessons that biological evolution has learned over the last few billion years.

We have been more concerned with efficiency and correctness than with robustness. This is sensible for developing mission-critical systems that have barely enough resources to perform their function. However, the rapid advance of microelectronics has alleviated the resource problem for most appli-

cations. Our increasing dependence on computational and communications infrastructure, and the development of ever more sophisticated attacks on that infrastructure, make it imperative that we turn our attention to robustness.

I am not advocating biomimetics; but observations of biological systems give us hints about how to incorporate principles of robustness into our engineering practice. Many of these principles are in direct conflict with the established practices of optimization and proofs of correctness.

As part of the continuing work to build artificially intelligent symbolic systems we have, incidentally, developed technological tools that can be used to support principles of robust design. For example, rather than thinking of backtracking as a method of organizing search we can employ it to increase the general applicability of components in a complex system that builds itself to meet certain constraints. I believe that we can pursue this new synthesis to get better hardware/software systems.

Bibliography

[1] Harold Abelson and Gerald Jay Sussman with Julie Sussman, *Structure and Interpretation of Computer Programs*, 2nd edition, MIT Press, ISBN 0-262-01553-0, (1996).

[2] Harold Abelson, Don Allen, Daniel Coore, Chris Hanson, George Homsy, Thomas F. Knight Jr., Radhika Nagpal, Erik Rauch, Gerald Jay Sussman, and Ron Weiss; "Amorphous Computing," in *Communications of the ACM*, **43**, 5, May 2000.

[3] *The ARRL Handbook for Radio Amateurs*, The American Radio Relay League, Newington, CT, USA (annually).

[4] Jonathan B.L. Bard; *Morphogenesis*, Cambridge University Press, (1990).

[5] Jacob Beal; *Generating Communications Systems Through Shared Context*, M.I.T. S.M. Thesis, also AI Technical Report 2002-002, January 2002.

[6] R. Bulirsch and J. Stoer; *Introduction to Numerical Analysis*, Springer-Verlag, (1991).

[7] E.M. del Pino and R.P. Elinson; "A novel developmental pattern for frogs: gastrulation produces an embryonic disk," in *Nature*, **306**, pp. 589-591, (1983).

[8] G.M. Edelman and J.A. Gally; "Degeneracy and complexity in biological systems," *Proc. Natl. Acad. Sci*, **98** pp. 13763–13768 (2001).

[9] M. D. Ernst, C. Kaplan, and C. Chambers. "Predicate Dispatching: A Unified Theory of Dispatch," In ECOOP'98. LNCS, vol. 1445. Springer, Berlin, 186-211 (1998).

[10] Robert Floyd; "Nondeterministic algorithms." in *JACM,* 14(4):636–644 (1967).

[11] Kenneth D. Forbus and Johan de Kleer; *Building Problem Solvers*, The MIT Press, (November 1993).

[12] Chris Hanson, `SOS` software: Scheme Object System, (1993).

[13] Carl E. Hewitt; "PLANNER: A language for proving theorems in robots." In *Proceedings of the International Joint Conference on Artificial Intelligence,* pp. 295–301 (1969).

[14] Paul Horowitz and Winfield Hill; *The Art of Electronics*, Cambridge University Press.

[15] Richard Kelsey, William Clinger, and Jonathan Rees (editors), *Revised*^{5} *Report on the Algorithmic Language Scheme*, (1998).

[16] Simon Kirby; *Language evolution without natural selection: From vocabulary to syntax in a population of learners.*, Edinburgh Occasional Paper in Linguistics EOPL-98-1, University of Edinburgh Department of Linguistics (1998).

[17] Marc W. Kirschner, John C. Gerhart; *The Plausibility of Life: Resolving Darwin's Dilemma*, New Haven: Yale University Press, ISBN 0-300-10865-6 (2005).

[18] Gregor Kiczales, `tinyCLOS` software: Kernelized CLOS, with a metaobject protocol, (1992).

[19] Harvey Lodish, Arnold Berk, S Lawrence Zipursky, Paul Matsudaira, David Baltimore, and James E Darnell; *Molecular Cell Biology*, (4th ed.), New York: W. H. Freeman & Co., ISBN 0-7167-3706-X (1999).

[20] David Allen McAllester and Jeffrey Mark Siskind; `SCREAMER` software.

[21] John McCarthy; "A basis for a mathematical theory of computation," in P. Braffort and D. Hirshberg, editors, *Computer Programming and Formal Systems*, pages 33–70, North-Holland, (1963).

[22] Paul Penfield Jr.; *MARTHA User's Manual*, Massachusetts Institute of Technology, Research Laboratory of Electronics, Electrodynamics Memorandum No. 6; (1970).

[23] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling; "Richardson Extrapolation and the Bulirsch-Stoer Method," in *Numerical Recipes in C: The Art of Scientific Computing*, 2nd ed. Cambridge University Press, pp. 718-725, (1992).

[24] Guido van Rossum, and Fred L. Drake, Jr. (Editor); *The Python Language Reference Manual*, Network Theory Ltd, ISBN 0954161785, (September 2003).

[25] http://www.python.org/dev/peps/pep-0020/

[26] Richard Matthew Stallman; *EMACS: The Extensible, Customizable, Self-Documenting Display Editor*, Massachusetts Institute of Technology Artificial Intelligence Laboratory Memo, AIM-519A (March 1981).

[27] Richard Matthew Stallman and Gerald Jay Sussman; "Forward Reasoning and Dependency-Directed Backtracking in a System for Computer-Aided Circuit Analysis," in *Artificial Intelligence*, **9**, pp 135–196, (1977).

[28] Guy L. Steele Jr.; Common Lisp the language, The Digital Equipment Corporation, (1990).

[29] Guy L. Steele Jr.; *The Definition and Implementation of a Computer Programming Language Based on Constraints*, MIT PhD Thesis, MIT Artificial Intelligence Laboratory Technical Report 595, (August 1980).

[30] Gerald Jay Sussman and Jack Wisdom with Meinhard E. Mayer, *Structure and Interpretation of Classical Mechanics*, MIT Press, ISBN 0-262-019455-4, (2001).

[31] *The TTL Data Book for Design Engineers*, by the Engineering Staff of Texas Instruments Incorporated, Semiconductor Group.

[32] Stephen A. Ward and Robert H. Halstead Jr.; *Computation Structures*, MIT Press, ISBN 0-262-23139-5, (1990).

[33] Daniel J. Weitzner, Hal Abelson, Tim Berners-Lee, Chris Hanson, Jim Hendler, Lalana Kagal, Deborah McGuinness, Gerald Jay Sussman, and K. Krasnow Waterman; *Transparent Accountable Data Mining: New Strategies for Privacy Protection*, MIT CSAIL Technical Report MIT-CSAIL-TR-2006-007 (27 January 2006).

[34] Lewis Wolpert, Rosa Beddington, Thomas Jessell, Peter Lawrence, Elliot Meyerowitz, and Jim Smith; *Principles of Development* (2nd ed.), Oxford University Press, ISBN-10: 0-19-924939-3, (2001).