October 24

2011

A Case for Sandia Investment in Complex Adaptive Systems Science and Technology

This white paper makes a case for Sandia National Laboratories investments in complex adaptive systems science and technology (S&T) -- investments that could enable higher-value-added and more-robustly-engineered solutions to challenges of importance to Sandia's energy and national security mission and to the nation.

Curtis Johnson
George Backus
Theresa Brown
Rich Colbaugh
Katherine Jones
Jeff Tsao

Cover Figure: Nature and engineering evolve similar solutions by means of adaptive processes.

Executive Summary

This white paper makes a case for Sandia National Laboratories investments in complex adaptive systems science and technology – investments that could enable higher-value-added and more-robustly-engineered solutions to challenges of importance to Sandia's energy and national security mission and to the nation. Our case, briefly, is the following:

- Complex adaptive systems are ubiquitous in Sandia's energy and national security mission areas. We often ignore the adaptive complexity of these systems by narrowing our "aperture of concern" to systems or subsystems with a limited range of function exposed to a limited range of environments over limited periods of time. But by widening our aperture of concern we could increase our impact considerably.

- To do so, the science and technology of complex adaptive systems must mature considerably. Despite an explosion of interest outside of Sandia, however, that science and technology is still in its youth. What has been missing is contact with real (rather than model) systems and real domain-area detail.

- With its center-of-gravity as an engineering laboratory, Sandia's has made considerable progress applying existing science and technology to real complex adaptive systems. It has focused much less, however, on advancing the science and technology itself. But its close contact with real systems and real domain-area detail represents a powerful strength with which to help complex adaptive systems science and technology mature.

- Sandia is thus both a prime beneficiary of, as well as potentially a prime contributor to, complex adaptive systems science and technology.

Complex Adaptive Systems Engineering

Complex Adaptive Systems S&T

Building a productive program in complex adaptive systems science and technology at Sandia will not be trivial, but a credible path can be envisioned: in the short run, continue to apply existing science and technology to real domain-area complex adaptive systems; in the medium run, jump-start the creation of new science and technology capability through Sandia's Laboratory Directed Research and Development program; and in the long run, inculcate an awareness at the Department of Energy of the importance of supporting complex adaptive systems science through its Office of Science.
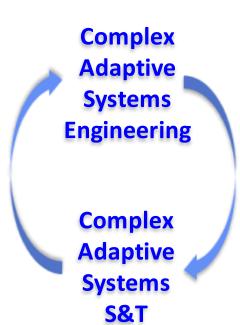
Table of Contents

0 Background

The intent of this white paper[1] is to propose a path for future Sandia National Laboratories (Sandia) investments in complex adaptive systems science and technology (S&T) – investments that could enable higher value and more robustly engineered solutions to a wide range of challenges of importance to Sandia's energy and national security mission and to the nation.

The white paper draws on interviews and feedback sessions with numerous Sandians[2] and non-Sandians,[3] as well as on key internal Sandia background documents.[4,5,6,7] These interviews and documents were invaluable to us as we converged on the point of view set forth in this white paper.[8] However, the white paper does *not* represent the least common denominator of those interviews and documents; rather, it represents a specific and somewhat unconventional point of view which was sharpened through those interviews and documents.

That point of view focuses on adaptive complexity: first, because it is ubiquitous in Sandia's mission areas (e.g., it is present in any adversarial environment and indeed any environment in which humans are involved); and second, because it causes system vulnerabilities and strengths that are critical for those mission areas but are often overlooked or underappreciated. Failure to address these vulnerabilities will allow adversaries to be always a step ahead. Capitalizing on these strengths will enable lasting solutions resistant to accident and attack.[9]

This white paper is intended to be accessible to a wide audience: to those inside and outside of Sandia, to the applications and "academic science" communities, and to those more and less knowledgeable about complex systems. To accommodate this range of audience, and because our point of view will not be familiar even to some knowledgeable about complex systems, we have adopted a pedagogical (intended to teach) rather than rhetorical (intended to persuade) style. This style is reflected in both the content and organization of the paper. The organization is Socratic, around a series of questions, which serve as the titles of the sections, and answers, which serve as the titles of the sub-sections. Thus, the table of contents serves a readable summary of the high-level technical content.

---

[1] This white paper was commissioned by Charles Barbour, Director of Sandia's Physical, Chemical and Nano Sciences Center 1100, and Director of the Enabling Capabilities Program Area within Sandia's ECIS (Energy, Climate and Infrastructure Surety) SMU (strategic management unit). We are especially grateful for cheerleading and moral support from Rush Robinett (Senior Manager of Grid, Modernization and Military Energy Systems Group 6110) and Russ Skocypec (Senior Manager of Systems Studies Group 240).

[2] Arlo Ames, Heidi Ammerlahn, Arnold Baker, Walt Beyeler, Rebecca Fang, Stan Fraley, John Ganter, Pablo Garcia, Bob Glass, Richard Griffith, Steve Hatch, Phil Heermann, Howard Hirano, Kevin Horn, Dan Horschel, Mark Ivey, Phil Kegelmeyer, David Keese, Steve Kleban, Mark Ladd, Kiran Lakkaraju, Gary Laughlin, Dawn Manley, John Mitchiner, Thomas Moore, Alan Nanco, Len Napolitano, Lori Parrott, Cindy Phillips, Mark D. Rintoul, Rush Robinett, Jonathan Rogers, Lani Sanders, Kent Schubert, Wendy Shaneyfelt, Russ Skocypec, David Stracuzzi, Steve Tucker, Tim Trucano, Jessica Westbrook, David Womble, and Tommy Woodall.

[3] Yaneer Bar-Yam (New England Complex System Institute), Cris Moore (University of New Mexico and Santa Fe Institute), Mark Newman (University of Michigan), and Geoff West (Santa Fe Institute).

[4] R.J. Glass, A.L. Ames, W.A. Stubblefield, S.H. Conrad, S.L. Maffitt, L.A. Malczynski, D.G. Wilson, J.J. Carlson, G.A. Backus, M.A. Ehlen, K.B. Vanderveen, D. Engi, "Sandia National Laboratories A Roadmap for the Complex Adaptive Systems of Systems (CASoS) Engineering Initiative," Sandia National Laboratories Report SAND 2008-4651 (September, 2008).

[5] A.L. Ames, R.J. Glass, T.J. Brown, E.B. Stechel, S.M. Deland, R.A. Suppona, P.D. Finley, T. Moore, S.H. Conrad, "A Roadmap for the Enterprise Security Engineering Initiative," unpublished.

[6] R. Skocypec, "Systems-of-Systems in National Security," presentation at IEEE Conference on System of Systems Engineering (Albuquerque, June 2, 2009).

[7] "Complex Adaptive Systems Engineering at Sandia: First Steps in Developing a Long-Term Strategy," white paper circa 2004 available upon request from Philip Heermann

(Senior Manager of Intelligent Systems, Robotics and Cybernetics Group 6530).

[8] But, of course, we absolve others of responsibility for deficiencies or errors that remain.

[9] Adaptation is nature's means of producing new capability for an uncertain and changing world. It is a blind process that improves performance incrementally through endless repetition – just the kind of approach taken by both biology and engineering. Thus, adaptive complexity offers one of our best hopes for a quantum leap in the intelligence of machines. Such a leap could, for example, give a massively parallel computer network a fighting chance against the dull-witted but persistent hacker who can consistently defeat it today by executing a series of trivial adaptations against its largely static and insensitive target.

1 What are complex adaptive systems?

A. *Systems whose structures modify to enable success in their environments;*

There is not yet a universally accepted definition for a complex adaptive system. Some definitions emphasize system structure (e.g., composed of many interacting and self-organizing parts), and some emphasize system behavior (e.g., emergent and difficult to predict from an understanding of the parts). In fact, there is a single process that both creates functional structure and enables emergence and other system behaviors: adaptation of system function to a changing environment. Adaptation entrains structure and behavior to each other and to the environment.

As adaptation drives both structure and behavior, it is the basis of our definition: *A system whose structure modifies to enable success in its environment.* In other words, a complex adaptive system displays adaptive behavior due to accumulated modification in the past *and* is able to adapt its structure and behavior to new challenges the environment may present to it in the future.

This definition is restrictive yet general.

It is restrictive in that adaption *must* involve structural modification. It is not enough for a system to vary as its environment changes – otherwise, a thermostat might qualify. It is not enough for a system to vary in an unpredictable or counterintuitive manner due to elaborate internal feedback loops – otherwise a system of weirdly coupled oscillators might also qualify.

It is general in that it places no restrictions other than that the system adapt, and that the adaptation involve structure modification. A system might qualify as a complex adaptive system regardless of: whether its designers are intelligent or dumb; whether it even has a designer; whether the adaptation is successful or not; or whether the mathematics required to describe the system are analytic, discrete, or purely algorithmic.

B. *That can be either biological or engineered;*

Perhaps most importantly, a system might meet our definition regardless of the raw material used to construct the system or of the mechanism for adaptation.

The system could be biological: adapting, without intent or foresight via random variation and the differential survival and propagation of genotypes. As illustrated by the tree on the left side of Figure 1, natural selection hones the structure and behavior of a biological system to enable the system to thrive under normal environmental variability: it modifies the system to better fit the system's function to its environment. Alternatively, the system could be engineered: adapting by design and artificial selection, as illustrated by the city of Paris on the right side of Figure 1.

Our definition emphasizes the one engine, adaptation, which creates the common properties in these two types of systems. While there are significant differences in the adaptive processes of biological and engineered systems, we argue that the analogy is tighter than is generally appreciated and is therefore underexploited.

First, like biological systems, engineered systems experience successes and failures under varying environmental conditions and undergo design changes based on those successes and failures. Breakdowns are repaired and improvements made, all to better fit the system's function to its environment, often to out-compete other systems occupying an overlapping niche in their common environment. So long as there is differential survival and propagation (i.e., the reduction and elimination of less successful candidates and the corresponding survival and growth of more successful ones) it matters not whether the mechanism is the natural selection of biology or the artificial selection of engineered systems.

Second, the analogy becomes tighter the grander the scale.[10] On modest scales, engineers can carefully and intelligently design and implement changes with the whole system in mind. Such changes are not at all like nature's random mutations and myopic adaptations. But over time an engineered system often grows beyond the capacity of any single engineer to understand or control and consists of an unintended patchwork of legacy and replacement components. On this grander scale, wholesale and omniscient redesign or replacement becomes impractical. The future of the system's structure necessarily falls into multiple hands, each having only local knowledge and local control, and therefore acquiring local interests.[11] Individual components

---

[10] Note that there is no discrete scale at which the system becomes adaptive. We cannot draw an absolute distinction any more that we can determine the exact step when chemical replicators transform into an organism or when a static algorithm has become a "learning" one. There is a continuum; and yet important qualitative differences do exist between distant points on it.

[11] This may begin, for example, with users choosing not to install patches or upgrades, or choosing to employ technology for unintended purposes or in unanticipated environments. Even when local hands report hierarchically to a single master, the potential for local interests and adaptation exists. Indeed, many of today's

Figure 1. Two examples of complex adaptive systems with deeply hierarchically modular structure: biological (left) and engineered (right). The similarity between the structures of the tree and of the city of Paris has two causes: both structures serve the function of transportation; and the two structures have evolved by similar processes of adaptive modification.

perform their functions in their local environment, experience successes and failures, and are modified accordingly: sometimes at the expense of other components, sometimes in cooperation, sometimes indifferently. Thus engineers are relegated to making local and myopic repairs and modifications and, as the city of Paris illustrates, engineered systems increasingly evolve just as biological systems do.

C. Whose hierarchical and modular order facilitate adaptation;

Systems that fit our definition share many qualities, perhaps the most important of which is that their structural organization is hierarchical and modular.

One important reason for hierarchical modularity was described in Herbert Simon's classic 1962 paper[12] introducing his now-famous parable[13] of the two

---

engineered systems are *designed* to come alive locally by enabling users to create their own new functionality (e.g., Matlab, Linux, Wiki, and many on-line games).

[12] H.A. Simon, "The Architecture of Complexity," Proceedings of the American Philosophical Society **106**, 467-482 (1962); and H.A. Simon, "The Sciences of the Artificial," 3rd Edition (MIT Press, 1996).

[13] There once were two watchmakers, named Hora and Tempus, who manufactured very fine watches. Both of them were highly regarded, and the phones in their workshops rang frequently. New customers were constantly calling them. However, Hora prospered while Tempus became poorer and poorer and finally lost his shop. What was the reason? The watches the men made consisted of about 1,000 parts each. Tempus had so constructed his that if he had one partially assembled and had to put it down – to answer the phone, say – it immediately fell to pieces and had to be reassembled from the elements. The better the customers liked his watches the more they phoned him and the more difficult it became for him to find enough uninterrupted time to

watchmakers Hora and Tempus. The moral of this parable is that it is much more probable that a system that is hierarchically modular (rather than an undifferentiated whole) can be constructed to completion in the face of inevitable mistakes or interruptions, and this difference in probability increases sharply with additional complexity.

Perhaps a more important reason for hierarchical modularity is that a system thus constructed is vastly easier to modify and enrich in structure and functionality as it evolves from one generation to the next. Components may be modified without breaking the system so long as the effects of those modifications are largely contained within a layer or module and obey the protocols governing traffic between the layers and modules of the system. Modification may enable groupings of modules into a new layer within the hierarchy (e.g., a market forms); while the subdivision of components may also spawn new layers (e.g., separating the functions of processing and memory in computing). And the reverse may also occur: whole layers and layer types may be replaced, destroyed, or discarded, if, for example, their cost exceeds their benefit.

The result is an explosion of adaptive variability accessible through incremental changes to individual

---

finish a watch. The watches Hora handled were no less complex than those of Tempus, but he had designed them so that he could put together sub-assemblies of about ten elements each. Ten of these sub-assemblies, again, could be put together into a larger sub-assembly and a system of ten of the latter constituted the whole watch. Hence, when Hora had to put down a partly assembled watch in order to answer the phone, he lost only a small part of his work, and he assembled his watches in only a fraction of the man-hours it took Tempus.

layers and modules. Because layers and modules become semi-independent and focused on their own subfunctions, the system comes alive internally and adaptation "goes local." Each layer and module adapts to its internal environment within the system just as the whole system adapts to *its* external environment. Thus adaptation drives modular hierarchy, which in turn accelerates adaptation.

This is true for both biological and engineered systems. In biological systems, the process of random variation and natural selection leads to layers and modules, e.g., genes, cells, organs, organisms, and groups of organisms. In engineered systems, the process of quasi-intelligent design and artificial selection also leads to layers and modules, e.g., transistors, integrated circuits, printed-circuit boards, software, computers, networks, and the Internet.

Thus, both biological and engineered systems adapt myopically through hierarchically modular structures. Through cumulative, incremental, and local adaptation they develop similar functional "fitness" relationships with their environments, both at the level of the system and at the level of any layer or module permitted to adapt with sufficient independence.

D. And whose goal-directed adaptations create non-random substructures.

Not only does hierarchically modular structure enable adaptation, adaptation in turn refines that hierarchically modular structure in specific ways. In particular, the various modules and links between modules, even within layers, become non-randomly different. The result is fundamentally different from a physical system such as a molecular gas, which is hierarchically modular but whose every layer[14] is composed of identical, randomly distributed entities.

Instead, the adaptive process is non-randomly selective. It accentuates whatever individual (and perhaps initially random and small) differences there might be amongst modules in a direction that best enables the system or subsystem to accomplish its function. In other words, the adaptive process is "teleonomic": it is directed towards the goal of improving how the system interacts functionally with its environment.

Indeed, this goal-directedness occurs for every semi-independent element at any level of the system. If we were to indulge in the anthropomorphic shorthand of biologists and speak of the individual system elements as if they were conscious, we would say that each element is competitive and entrepreneurial, doing business on the fly with the best local "providers" and "customers" available to it, all to further the success of its local function and to further the larger aim of the system to improve *its* function.

The result is differentiation and specialization, with elements out-competing each other for various functional niches by employing different modes[15] of achieving those functions. Thumbs are not the same as fingers, even though they are both modules at the same level in the hierarchy of the human body. Functional and goal-directed adaptation creates non-random substructure within and at all levels of the overall hierarchically modular system structure. It is the basis for the "organized complexity"[16] associated with complex adaptive systems, as opposed to the "disorganized complexity" associated with non-adaptive systems.

However, just because a system has non-random and organized substructure does not mean that a particular substructure is predictable or inevitable. Any goal (or end) can be accomplished through a number of functional structures (or means). Which functional structure is actually used depends on the details of the evolutionary path taken and on historical accident. Likewise, functional structures can be co-opted to create new functions and goals. Evolution is a tinkerer, making use of whatever components are at its disposal to improve existing – or to produce new – function.[17] Thumbs enhanced our ability to grasp objects and subsequently led to the new functions of gesturing and tool-making.

Thus, just as a particular (hierarchically modular) structure facilitates goal-directed adaptation, goal-directed adaptation in turn creates and superimposes non-random and functional substructures onto the hierarchically modular structure. Complex adaptive systems are not molecular gases or sandpiles; they are goal-seeking systems whose structures are adaptationally "alive."

---

[14] The various layers being: the gas itself; the molecules that the gas is composed of; the atoms that the molecules are composed of; the protons, neutrons and electrons that the atoms are composed of; and so on.

[15] Through, e.g., the links through which they trade information and material with other parts of the system.

[16] W. Weaver, "Science and Complexity," American Scientist **36**, 536 (1948).

[17] F. Jacob, "Evolution and Tinkering," Science **196**, 1161-1166 (1977).

2 Why should Sandia care about complex adaptive systems?

A. Because they are ubiquitous in biological, engineered, and human sociotechnical systems;

As discussed above, our definition of complex adaptive systems places no restriction on the raw material used to construct the system or on the process used to achieve adaptation. It applies equally to biological life and to engineered systems. On the time scale of organismal life cycles for biology, and of product cycles for engineering, hierarchically modular systems that have adapted over many generations to fit their environments are ubiquitous. Indeed, as has been said of networks, "once you begin to study them it is difficult not to see them everywhere."[18]

Perhaps most importantly, complex adaptive systems are also ubiquitous in sociotechnical systems at the intersection of biology and engineering. Humans have evolved to be adaptive organisms; their technologies are engineered; and because of the tight dependence of modern humans on their technologies, virtually every system that involves humans as producer, customer or operator is a sociotechnical, complex adaptive system.

B. And, most importantly, because they are ubiquitous in Sandia's energy and national security mission areas.

Put in this broad perspective, virtually all systems of interest to Sandia in its energy and national security mission areas are complex adaptive sociotechnical systems involving both humans and technology. Resource distribution systems have humans producing or extracting those resources at one end and consuming them at the other. Weapons systems have humans firing at other humans or their habitats. Surveillance and security systems often scan humans and must not be outwitted by them. Here, we briefly mention three examples of complex adaptive sociotechnical systems of past and future interest to Sandia.

A first example is digital information systems, which have spawned perhaps the greatest complexity and evolved the fastest of all sociotechnical systems.

The first computers had a layer or two of hardware whose "programs" were hard-wired on plugboards – the programmer literally configured the hardware. Many layers have since been added. A true software layer was added in the form of punchcards. In rapid succession, affordable desktop machines appeared with complex hierarchies of hardware and software layers; personal computers (PC's) were connected to peer machines and other devices in local networks and then to the Internet (layers upon layers); and wireless channels were added. Static content was replaced by dynamic content (interactive communication). Commerce was added. One-to-a-few communications were supplemented by one-to-many systems such as Facebook and Twitter. Spreadsheets replaced calculators and manual accounting books. Word processors replaced typewriters. This tremendous evolution in the hierarchical complexity of digital information systems has greatly enriched system functionality, and created and irrevocably altered entire industries.

Digital information systems are ubiquitous in military hardware, sensing and surveillance systems, intelligence collection and analysis, security systems, and infrastructure. Virtually all of these systems share information at some frequency with the global information grid and are vulnerable to attack. They are critical for national energy, economic and military purposes, and they have become the most complicated and difficult of infrastructures to defend.

A second example is the smart electricity grid, an emerging system construct which will necessarily be constructed from many modules organized into a hierarchy of layers. These layers include:[19] an electron layer (production, transmission and consumption of electrical power); a financial/economic layer; an information layer (including those information systems embedded in appliances); one or more social layers (including politics, popular opinion, consumer and corporate behavior[20]).

The physical and social scale of such a system is enormous, and its rich and complex functionality necessarily entails not only these multiple layers, but communication and power-transfer protocols across layers. Multiple actors, such as utilities, regulators, and consumers interact in the market with often incompatible goals (i.e., functions) and with potentially fragile physical and financial information flows.

Most importantly, implementing a smart electricity grid will involve a transition from centralized control

---

[18] S. Goyal, "Strategic Network Formation," presentation at GameNets 2009 (Istanbul, 2009).

[19] National Science and Technology Council, "A Policy Framework for the 21st Century Grid: Enabling Our Secure Energy Future" (June 2011).

[20] Some behavior might be aimed at maximizing profit; some might be aimed at other goals, including nefarious ones. And all behavior might to some extent employ market-gaming strategies of various kinds.

SMART GRID
A vision for the future — a network of integrated microgrids that can monitor and heal itself.

Smart appliances
Can shut off in response to frequency fluctuations.

Demand management
Use can be shifted to off-peak times to save money.

Solar panels

Offices

Houses

Processors
Execute special protection schemes in microseconds.

Sensors
Detect fluctuations and disturbances, and can signal for areas to be isolated.

Disturbance in the grid

Storage
Energy generated at off-peak times could be stored in batteries for later use.

Isolated microgrid

Wind farm

Generators
Energy from small generators and solar panels can reduce overall demand on the grid.

Industrial plant
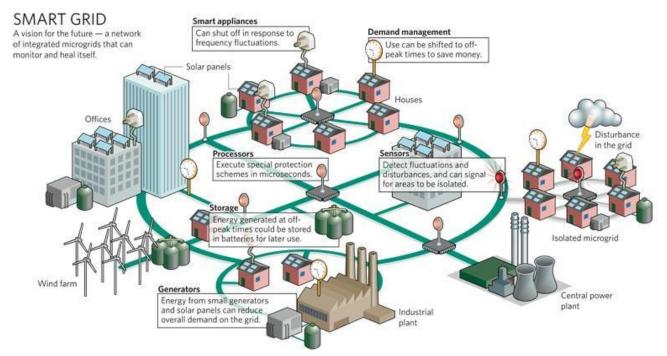
Central power plant

Figure 2. Schematic of one (or perhaps a few) level(s) of a hypothetical and highly hierarchically modular smart electricity grid [after E. Marris, "Energy: Upgrading the grid," Nature 454, 570-573 (2008)].

to a fully decentralized market process in which the inertia associated with the physical and human components of the system may prevent smooth and adiabatic progress from one set of stable protocols to another. The transition from a stable "before" to a stable "after" will produce complex time-sensitive and path-dependent trajectories.

A third example is a nuclear energy system.[21] Here, the system is an international community of nation states powered in part by nuclear energy. Because processes associated with the nuclear energy life cycle can be co-opted to create nuclear weapons, and nuclear weapons bring perceived local benefits to nation states, individual nation states have incentives to develop such weapons. However, because nuclear proliferation also brings global risks, it is of interest to design policies or behaviors on the part of a few cooperative nation states which would minimize proliferation.

Such a system is also necessarily constructed from a hierarchy of levels, spanning: the nuclear fuel and processing life-cycle itself; the nuclear power plants which touch that life cycle but also represent their own subsystem of operators, control procedures, guards and fences; and organizations (individuals, terrorist groups, nation states, international bodies)

with various motivations to limit or engage in nuclear proliferation and with existing but evolving non-nuclear interactions including sociopolitical (conflicts and alliances), economic (trade or economic dependency), and cultural (shared or unshared heritages). And such a system is especially difficult to treat because of the degree to which players must act in the face of uncertain knowledge and the range of timescales involved.

C. Of course, by narrowing its "aperture of concern" Sandia can limit itself to treating only non-complex subsystems or components of these systems;

As discussed above, it is possible to approach nearly every energy or national security system of interest to Sandia as a complex adaptive system. Whether one actually should do so, however, depends on the size of one's "aperture of concern" within one of four dimensions. We illustrate these dimensions using airport security as a specific example.

The first dimension is the purposes and outcomes of the project. The precision and accuracy of a physical detector is a narrow outcome. The scope can be *broadened* to include outcomes in other domains (uptime; throughput; cost; effect on airport businesses, customs, and air traffic control; broader economic impact; effect on TSA operations; privacy and civil liberties impacts; effect on cosmetic and shoe manufacturers, etc.). The scope can be *heightened* by embracing the ends for which the current goal is a

---

[21] A.L. Pregenzer, R.J. Glass, A. Ames, W.E. Beyeler, S. DeLand, and A.D. Williams, "A Systems Approach to Assessing Nonproliferation Strategies" (draft report).

means (e.g., by asking "why" with regard to the current objective). In this way, the goal rises from explosives detection to hazard-free passengers and luggage, and from there to the prevention of aviation terrorism while maintaining an economically viable air-transportation system. Expanding purposes and outcomes brings many new layers and modules into the system.

The second dimension is the range of *random* environmental variation the system is expected to be robust across. These environmental variations might include: explosives and precursors; chemical, electrical or magnetic interference; ambient temperature and humidity; passengers (e.g., handicaps and clothing); luggage and luggage contents; and operators (e.g., their intelligence, skills and attentiveness). This range can be most easily expressed by the variability that can be expected with a certain probability in a fixed time period (e.g., a hundred-years' flood). Certain aspects of environmental variability lead to errors and breakdowns in a given engineered system, which catalyze modifications. The system will be hard-wired for a certain range of variability and will (tacitly or expressly) rely on repair, upgrade or replacement for the remainder. Here, the organized complexity is the continual selection pressure that causes the system to adapt to its environment. Expansion of the aperture in this second dimension leads to flexible designs and systems that "learn" from experience.

The third dimension is closely related to the second. It is the range of *non-random* variability in the environment. While the second dimension assumes a disinterested environment, this third dimension accounts for "interested" environments and the coevolution of the system with them. By far the most frequent sources of non-random variability in national security systems are humans. Feedback from their own experience and other information sources lead humans to change behaviors. People adapt to the airport security system, changing their dress, their luggage contents and their behavior to achieve their own ends, which may be innocent (e.g., avoid embarrassment or delay) or nefarious (e.g., defeat detection). Expanding the aperture in the third dimension generally means abandoning a "human factors" approach in favor of an engineering approach that includes the human components within the boundaries of the system. In other words, static design can go only so far in anticipating the tactics of an adversary. At some threshold, a robust system must customize its reactions to an adversary, perhaps finding the means to shape the adversary's responses.

The fourth and last dimension is the extent to which the scope embraces the people, equipment, and processes that *produce* the "final" engineered system. A deceptively simple part is the product of a complex production system. It is easy to assume that contracts ensure that producers are serving their customers' interests. While suppliers and customers share many objectives and benefit from many of the same results, this alignment is never perfect. The manufacturers of security systems seek profit; TSA seeks credibility, reduced terror risk, and self-preservation; airports and airlines want passenger throughput. Risks increase dramatically if we add the possibility of malicious intent. If national security depends on the reliability of a part, it may be necessary to include the supply chain for the part in the project scope in order to address risks from manufacturing reliability to counterfeit parts or sabotage. The selection and training of operators and maintenance staff are also be part of this dimension.

D. But there is tremendous value-added for the "aperture of concern" to be as wide as possible;

On all four dimensions, the wider one opens the aperture to organized complexity the higher the potential value to Sandia's energy and national security mission. Any solution wrought in the broader context will be more comprehensive, robust and enduring. And, the more advanced the S&T of complex adaptive systems, the wider are the apertures that can be practically achieved. Thus, advances in complex adaptive systems S&T, particularly those that expand its scope along our four dimensions of organized complexity, are of great interest to Sandia and to its energy and national security mission.

E. As well as consistency with an inevitable shift in Sandia's research enterprise towards complex adaptive systems S&T.

Moreover, moving in this direction is consistent with a likely shift in the character of Sandia's S&T research enterprise over the coming decades.

First, Sandia's S&T center-of-gravity will likely shift from subsystems and components towards systems. The reason: global economic trends. As lower-value-added manufacturing of components and subsystems has moved to low-wage nations, the U.S. has increasingly shifted to higher-value-added systems integration. For example, the successes of three of today's most transformative consumer products are based clearly on system advances: Google's PageRank search algorithm, Facebook's tools for social network analysis and data mining, and Apple's hardware and software ecosystem for smart mobile devices such as the iPod, iPhone and iPad.

Second, even as this shift towards systems S&T gains momentum, the systems approaches that are being applied will likely increasingly take on the character of complex adaptive systems approaches. In part this is because the explosion in power of underlying computational and communication components has enabled systems of unprecedented connectivity and reach, whose understanding *requires* a complex systems approach. It is also, in part, because that same explosion in computational power *enables* complex adaptive systems approaches. Systems that were previously not, now are, amenable to such approaches; and systems can now expand to include crucial components and interactions which previously were swept artificially into a fixed "environment." In other words: on the one hand, systems increasingly contain components (including humans) which are smart, communicative and adaptive; while on the other hand, system designers and engineers increasingly have the computational power to understand and treat systems containing such smart, communicative and adaptive components.

Indeed, Sandia has taken small steps in this direction for more than a decade. In 1996, Sandia developed perhaps its first complex adaptive systems capabilities: its Aspen agent-based model, used to study both economic development and its potential disruption through terrorist acts.[22] Other early complex adaptive systems studies included the relationship between the North American Free Trade Agreement (NAFTA), border guard behavior, and the detection of explosives in the field by an array of mobile robots. More recently, Sandia's National Infrastructure Simulation and Analysis Center (NISAC),[23] making use of network and agent-based simulation tools, has studied a wide diversity of complex adaptive systems, including one which has directly informed national public health policy: pandemic-resistant community networks.[24] Other examples include evaluation of global supply chain risks due to disruptions within the supply chain or to infrastructures on which the chain depends, and evaluation of the potential impacts of hurricanes and earthquakes to support planning and preparedness activities.[25]

We recognize, however, that this shift will not be pain-free. Sandia's current center of gravity is clearly

in subsystems or components: improving them as a means to improve the performance of their larger systems. Sandia's S&T investments reflect this: only one of its six research foundations,[26] computer and information sciences, intersects in any significant way systems S&T, and only about 1/10 of Sandia's work is considered systems related[27]. Expanding systems-related S&T at the expense of subsystem and component S&T will, at the Sandia institutional level, be difficult. Indeed, a significant question will be how to position Sandia's physical sciences S&T so that it can continue to be relevant in a world where Sandia has made the shift in emphasis to complex adaptive systems S&T.

3 What are some unique properties of complex adaptive systems of importance to Sandia's mission?

As just discussed, complex adaptive systems are ubiquitous in Sandia's energy and national security mission areas. The reason is simple: both biology and engineering have converged on complex adaptive systems based on hierarchical modularity and adaptation at all levels in the hierarchy as the means to enable and maintain high levels of performance in unpredictably changing environments. Complex adaptive systems have an overwhelming advantage over non-adaptive systems, and are the reason for the richness and success of both biology and engineering. Neither pheromones nor TCP/IP would pervade their respective ecological niches without the benefit of generations of adaptation.

Along with hierarchical modularity and adaptation at all levels in the hierarchy, come unique properties. These properties are strengths that enable the success

---

[22] N. Basu, R. J. Pryor, T. Quint, and T. Arnold, "Aspen: A Microsimulation Model of the Economy," SAND96-2459. Albuquerque, NM: Sandia National Laboratories (1996).

[23] http://www.sandia.gov/nisac/.

[24] http://www.sandia.gov/CasosEngineering/pandemic_influenza.html.

[25] http://www.sandia.gov/nisac/analyses.html.

[26] Currently, these are: Bioscience, Computer and Information Sciences, Engineering Sciences, Materials Science & Technology, Microelectronics & Microsystems, and Pulsed Power.

[27] Although systems S&T is a small fraction of Sandia work, it has a long and distinguished history. Since its birth in 1949, when it was asked to take responsibility for the development and integration of the non-nuclear (ordnance) pieces of our nation's nuclear weapons, Sandia has necessarily taken a systems perspective on design and engineering. This perspective has included not just the interactions between components within the system, but interactions between the system and its environment. The strong/weak link and exclusion zone mechanism is an early and elegant example of systems design and engineering for the safety, security and reliability (surety) of nuclear weapons systems in unpredictable and sometimes hostile physical and social environments.

---

of complex adaptive systems. However, they also represent weaknesses: vulnerabilities that can be taken advantage of.

In this section, we discuss three of these properties, ones which are of special importance to Sandia's energy and national security mission areas. We do not discuss how one might mitigate the weaknesses associated with these properties, except to note here that such mitigation is a major motivation for the development of the deeper and richer complex adaptive systems science and technology discussed in subsequent sections of this white paper.

A. First, structure, function, and environment are mutually adapted;

A first property unique to complex adaptive systems is the mutual adaptedness of structure, function, and environment. Adaptation can be thought of as a process of incorporating environmental information into system structure and function. The system "knows" how to avoid environmental dangers, to take advantage of environmental opportunities, and to prepare for environmental variability that it has experienced in its recent past. This knowledge can be stored in its structure: for example, a camouflaged species literally wears a copy of its visual environment on its back – plain to see (if you are able).

Indeed, because structure, function, and environment are so highly interdependent, qualities of any one of the three can be inferred from the other two: function is revealed by what the system controls and optimizes in the face of normal environmental variation; recent environmental variability can be deduced from system tolerances; and unexplained, little-used structures may be latent controls, or "scars," developed for distant-past or infrequent environmental conditions.

Thus, systems that have evolved over time are peculiarly robust to their environment. They have experienced all of the normal variability of that environment and survived it: they have found ways to achieve necessary functions throughout the ups and downs.

On the one hand, this robustness to normal environmental variability is a tremendous strength. Organisms know how to store food, avoid extreme temperatures, and sleep, all in order to conserve energy for times when opportunities are greater. They know how to keep their eggs warm and to evict parasites. They know how to synthesize (expensively) only the vitamins they cannot do without and cannot acquire more cheaply from their environment. Animals that depend on unreliable sources of food or other necessities go extinct;[28] the ones that remain are robust.

In other words, systems develop stable interfaces to those aspects of their environment which are themselves inherently stable, and they create control systems to manage the variable aspects of their environment on which they depend. Normal perturbations tend to have only temporary effects on such systems. Swatting flies in the backyard only makes room in the ecology for younger flies – the population is robust to losses because it always produces many more young than the environment can support. The security advantages of such a system over a similar but non-adaptive system, no matter how well designed, are obvious.

On the other hand, underneath this robustness to normal environmental variability lies a weakness. Adaptive robustness is not absolute, but specific both to the environment and to the functions performed. The more finely tuned a system is, the more likely it will struggle to perform other functions or to succeed in new environments. Fragilities are therefore to be expected at and beyond the normal range of variability and in less critical functions.

B. Second, adaptive elements co-evolve;

A second property unique to complex adaptive systems is the coevolution of adaptive elements: adaptive systems and their elements modify their behaviors in response to and in anticipation of each others' behavior. The relationships between individual elements can take multiple forms over time or even at a given moment. As illustrated in figure 4, these forms include: quid pro quo, adversarial, cooperative, and competitive. Moreover, also as illustrated in figure 3, these intra-system relationships among elements occur at all levels of the system as well as between systems in the environment. Elements in the system become locally adaptive even as they are each part of larger groups and higher layers.

The competitive aspect of the relationship between system elements can be an especially strong driver of both system behavior and adaptation, with individual elements competing for information, resources and relationships.

On the one hand, this competition among diverse elements is a great strength, fueling differentiation and specialization. If, for example, one module (a heart for pumping blood) in a biological organism is more efficient than another at providing a particular

[28] It has been speculated that locusts have evolved prime-number emergence cycles so that no predator with a more frequent cycle can entrain with it.
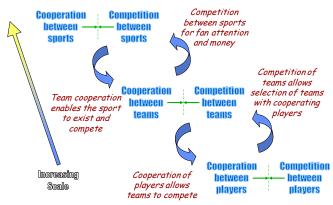
Figure 3. Example of how different layers of a complex adaptive system can "come alive," with both cooperative and competitive relationships between system elements catalyzing further adaptation [after "Complex Adaptive Systems Engineering at Sandia: First Steps in Developing a Long-Term Strategy," white paper circa 2004 available upon request from Philip Heermann (Senior Manager of Intelligent Systems, Robotics and Cybernetics Group 6530)].

function, then over biological evolutionary time it may become the one type of module that provides that function. Or if, for example, one component in a smart electrical grid is more efficient than another at power switching, or if one communication protocol, like TCP/IP, provides more functionality than another, over engineering evolutionary time these may outcompete alternatives globally or in a given competitive niche. In other words, modules specialize and find niches in which they have superior performance, creating ever more sophisticated and finely tuned system performance.

On the other hand, intra-system competition can be a weakness. Local and global selection pressures, though they mostly overlap, can also diverge.[29] The resulting system can be very different from what an intelligent designer with a perfect view of the whole system and its environment would make. Some adaptations will inevitably be in tension or competition with each other, some even maladaptive for the system as a whole (e.g., cancer). These maladaptations represent vulnerabilities that an adversary could exploit.

Moreover, the interactions of elements, like those characteristic of games in general, tend toward stable states. Consistent, robust strategies generally emerge. However, these strategies are robust only to their own recently experienced environments, in this case the recently experienced strategies of other elements. New strategies and strategies invading from other environments are often disruptive.

Finally, interactions do not always tend toward stable states, but can themselves evolve dynamically. New variants of an original vulnerability may arise through adaptation (just as bacteria evolve to keep pace with antibiotic treatment); followed by a fixing of the new vulnerability; and followed in turn by a new variant vulnerability.

C. Third, simple protocols govern traffic and exchange among modules and layers;

A third property of complex adaptive systems is the existence of simple but pervasive protocols which define the nature, path, and format of traffic and exchange among modules and layers.

Such protocols are not required in systems that are simple and can tolerate specialized, "hard-wired" paths and rules unique to the two elements they join. As a system grows in size and complexity, however, so also does the traffic within the system. More languages are spoken, more types of goods are moved in a growing range of quantities across a growing range of distances through an increasing number of intermediaries.

Such protocols, then, are necessary in complex adaptive systems composed of semi-independent elements which contribute to, and benefit from, the system as a whole. Individual elements must be able to interact with other elements, even with "stranger" elements previously unknown to it. Complex adaptive systems fall apart if they are unable to solve this interaction problem by developing common protocols: common languages, transport container standards, and means of exchanging goods and value (e.g., markets).

On the one hand, the existence of simple protocols is a great strength of complex adaptive systems. Cooperation amongst modules on behalf of the system generally cannot evolve without protocols to prevent the worst kinds of "cheating" among modules. Protocols are what prevent a complex adaptive system from the tragedy of the commons or from being populated by defectors.

On the other hand, the existence of simple protocols brings weaknesses and vulnerabilities. Protocols reveal what is trafficked and with what priorities (e.g., the tradeoffs among speed, capacity, efficiency, flexibility, and reliability), and are highly vulnerable to exploitation.[30,31] Protocol specifications are often

---

[29] Perhaps the most important challenge facing evolutionary biology today is understanding such multi-level selection: the interaction of selection at multiple layers.

[30] An example of protocol vulnerability in software systems is a buffer overflow. Software typically includes instructions for writing partial results to memory. Hardware and firmware layers determine the sizes and locations of the temporary memory buffers where these results are stored. It is not unusual for the data to be too

public and, when not, are easily reverse engineered. The protocol describes in detail acceptable inputs to a module. Those who wish to penetrate that module know precisely the shape and color of the Trojan Horse that guards will permit through the gates. And even if the protocol vulnerability is known, it can be exceptionally hard to develop a new protocol that preserves system performance and to win "agreement" for it. A globally "accepted" protocol is perhaps the most stable element of a complex adaptive system and therefore the most difficult to fix.

D. Taken together, these three properties create unique challenges and opportunities.

The dilemma and challenges are now clear. Complex adaptive systems evolve based on hierarchical modularity and adaptation at all levels in the hierarchy. This architecture has overwhelming strengths (benefits), but necessarily entails weaknesses (vulnerabilities).

Better engineering can reduce and mitigate vulnerabilities but cannot eliminate them. In fact, the controls imposed to mitigate small frequent upsets necessarily generate fragilities associated with other structures (so-called "robust yet fragile" or RYF behavior[32]). These are challenges for all engineering, but they are especially important for engineering systems that must be extremely reliable and that could be subject to persistent, intelligent attack: precisely the energy and national security systems discussed in section 2.

On the flip side of the coin, a system with semi-independent parts monitoring each other has the potential for efficient, robust, distributed, layered defenses. When the adaptive nature of the system is ignored or even resisted, "fixes" are likely to be fragile, fleeting, and particularly vulnerable to an intelligent adversary. But where adaptation can be leveraged, "solutions" become robust and enduring against

accidental and intentional upsets, and difficult for an adversary to pick apart.[33]

If, through advances in complex adaptive systems S&T, such solutions could be found, they would have dramatic impact on Sandia's mission, in domain areas such as energy, cyber, critical infrastructures, and security and weapons systems.

4   What is the current state of complex adaptive systems S&T?

A. It is young, but maturing as it increasingly connects to real-world observation and experiment;

Complexity science, per se, arose from the investigation of disorganized complexity in physics and chemistry, beginning in the 1970's[34] (perhaps even earlier[35]) and accelerating in the 1990's[36] and 2000's.[37]   Much progress has been made,[38] particularly through the study of idealized model systems. Tremendous advances have been made in the mathematical and computational tools for analyzing, and in our theoretical understanding of the behavior of, these model systems. However, it is only when fuller contact is made with real-world observation and experiment that a science can really "take off."[39] Until then, one considers a field to be in the early stages of development, and this is where we consider complex adaptive systems S&T to be right now.

---

large for the target memory. The excess data overflows the buffer and writes over some other portion of memory. Many malware attacks perpetrate buffer overflows to achieve their goals.

[31] A fctional example of a protocol exploitation was discussed by Vernor Vinge in his novel "A Deepness in the Sky" (Tor Books, 1999), in which the hero wins by exploiting an obscure aspect of one layer in a giant communication stack whose protocol he helped write 40 years before.

[32] J.M. Carlson, J. Doyle, "Complexity and robustness," Proc. Natl. Acad. Sci. USA **99**, 2538-2545 (2002); R.A. LaViolette, K. Glass, R. Colbaugh, "Deep information from limited observation of robust yet fragile systems," Physica **A 388**, 3283-3287 (2009).

[33] Y. Bar-Yam, "Engineering Complex Systems: A New Paradigm," Chapter 2 in Complex Engineered Systems: Science Meets Technology, D. Braha, A.A. Minai, Y. Bar-Yam, Eds. (Springer, 2006); Y. Bar-Yam, "Making Things Work: Solving Complex Problems in a Complex World" (NECSI Knowledge Press, 2004).

[34] P.W. Anderson, "More is different," Science 177, 393-396 (1972).

[35] See, e.g., B. Castellani's "A History of Complexity Science" map: http://www.personal.kent.edu/~bcastel3/complex_map.html.

[36] G.W. Flake, "The Computational Beauty of Nature" (MIT Press, 1998).

[37] M. Mitchell, Complexity: A Guided Tour (Oxford University Press, 2009).

[38] M. E. J. Newman, "Complex systems: A survey," Am. J. Phys. **79**, 800-810 (2011).

[39] As did the physical sciences in the 17th century, coincident with the Royal Society's adoption for its motto the phrase "nullius in verba" (take nobody's word for it) to emphasize the growing importance of observation and experimentation.

In the past decade or so, however, there have been glimmers of progress in making such contact with real-world observation and experiment. The recent availability of real-world data sets from systems as diverse as the Internet and biology have, e.g., arguably[40] been responsible for an unprecedented explosion of work in network science. With the accelerating ubiquity of sensors, computation, and communication, there will be almost certainly an accelerating availability of such data sets from real-world systems. There is nothing like good data to stimulate and test new theories. As a consequence, we see complex adaptive systems S&T poised to advance in the coming decade.

Two major difficulties must be surmounted, however, in order to facilitate contact with real-world observation and experiment. These two difficulties are illustrated in the generalized flow diagram in figure 5: the first, having to do with theory, lies on the left side of the diagram; the second, having to do with experiment, lies on the right side of the diagram.

B. But fundamental new approaches are needed to model systems with real-world adaptive complexity;

The first difficulty is that current approaches to systems modeling are unable to handle real-world adaptive complexity. Complex adaptive systems are simply not amenable to the purely reductionist approaches that have been so successful in physical sciences in the past,[41] and that understandably continue to bias current research agendas. The reason: reductionism must either retain atomic elements in all their gory detail, or else average and aggregate their effects with approaches like mean-field.

On the one hand, detailed reductionism, the full reduction of a complex system to its parts, is computationally impossible and self-defeating. The data usually do not exist in the detail necessary to construct the model, and even where data exist, they do not usually connect to the system phenomena of greatest interest.[42] This is not to say that insights cannot be gleaned from detailed reductionism. The tools of systems engineering (e.g., waterfall charts and wiring diagrams), high-performance computing (e.g., scaling algorithms, massively parallel agent-based models and Monte Carlo simulations) have been highly successful at providing insight, much counterintuitive, into the potential effects of real and hypothesized environmental perturbations across time and space. Nevertheless, it is difficult to identify and track the occasional large impacts of seemingly small and irrelevant environmental perturbations, leading some practitioners to a "you may never know" approach to complex adaptive systems modeling, or to non-quantitative descriptive models or case studies.

On the other hand, mean-field reductionism, the use of averaging techniques to aggregate randomly behaving constituent elements into simplified mean-field-behaving classes of elements, can reveal higher level order in messy phenomena (e.g., Brownian motion). Properly applied, these techniques have been tremendously useful in complex adaptive systems. However, they can also abstract away the structures that are the most important to preserve: the organized functional complexity of the system. As discussed in section 1, in a hierarchically modular system, lower-level "particulate" elements and events are neither identical nor randomly different, but are non-randomly different. Particularly in real-world systems, both function and diversity evolve through an accumulation of environmental interactions over time, many of which are simply historical accidents that are difficult to abstract effectively, but have significant effect on performance and future adaptation.

Thus, neither of these reductionist approaches has been able to handle real-world adaptive complexity, and the challenge is to develop fundamental new approaches that *can* handle such real-world adaptive complexity.

---

[40] A.L. Barabasi, closing comments at NetSci 2010 (Boston, May 20-14, 2010).

[41] For example, Lavoisier's breakthroughs in chemistry exemplify the tremendous explanatory power achieved by this method. Previously inexplicable macroscopic phenomena such as fire and crystallization were understood by microscopic probing of ever-deeper layers and ever-smaller spatial scales.

[42] Where detailed data exist is often where the system is already best understood and, more importantly, where the system is likely to be robust to environmental variability. The myopic focus on recent events and functional goals of complex systems causes them to evolve "control systems" that enable consistent performance in the face of normal environmental variability. One therefore expects robust controls to be in place wherever frequent variability endangers functional outcomes. Detailed modeling efforts often fail to take advantage of this fact and, in effect, provide detailed models of already solved problems: modeling both the environmental variability and the highly effective system that negates the effects of that variability. For these reasons, detailed models often parrot back what we know (while inevitably introducing new sources of error) far more than they reveal what we don't know.

THEORY | EXPERIMENT

Abstracted System ← Insight → Actual System

Analysis — Network (Fine Grained) | Mean Field (Coarse Grained) — Prediction — Insight — Observation — Test — Artificial (Experiment) | Natural (Field Study)

Re-Designed Abstracted System → Insight → Re-Engineered Actual System
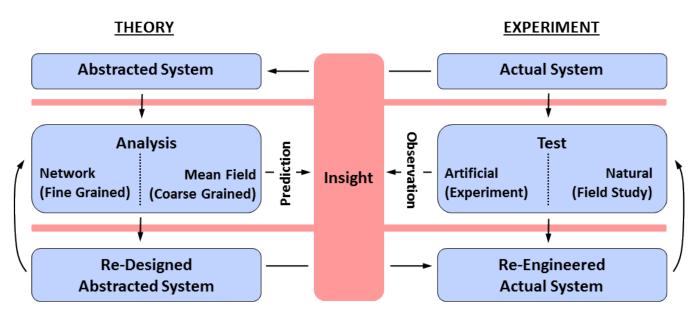
Figure 4. Generalized flow diagram associated with the "actual" (right column) and "abstracted" (left column) aspects of a system. We start in the right "experiment" column with some kind of actual system. If theory is not involved, then the system enters a test and re-engineering loop that ends when the system's function (how it interacts with its environment) is satisfactory. If theory *is* involved, then the system must be abstracted in some way, so that it can be analyzed. Then the abstracted system enters an analysis and re-design loop that again ends when the abstracted system's function (how it is expected to interact with its environment) is satisfactory. Finally, the re-designed abstracted system is un-abstracted and instantiated in a re-engineered actual system.

C. And new observational and experimental platforms must be developed for generating data on real-world adaptive complexity.

The second difficulty is that real-world complex adaptive systems are, by definition, much more complex than model systems, and usually engineered at a scale that prohibits any single designer from being "in charge." As a consequence: measurements are difficult to make; when they are made they are not necessarily of the quantities that would best elucidate system structure or function; and even if the right quantities are measured they are not usually made under artificially controlled (but rather under natural) conditions.

The challenge is thus to develop technologies that enable measurements that: are easy to make, and/or are of quantities that best elucidate system structure or function, and/or are made under artificially controlled, not just natural, conditions. An example of the first is the development of GPS (global positioning satellite) enabled mobile phones, which has enabled measurement of the spatial position of humans embedded in networks of various kinds. An example of the second is the development of fMRI (functional magnetic resonance imaging), which has enabled measurement of structure/behavior relations in the brains of humans – a critical component of many complex adaptive networks. An example of the third

is the use of online social networks as platforms for artificially controlled experiments in social interactions.

As illustrated on the right half of the figure 4 flow chart, the challenge here is to start with an actual real-world system, then figuring how best to re-engineer it so that it can be a better platform for comparisons between theory and hence from which to extract scientific insight. In other words, to bring to bear on complex adaptive systems the observational and experimental machinery that has proven so beneficial in the physical sciences.

That said, field studies are also essential. Skinner box experiments and studies of captive animals, while necessary, are insufficient to reveal – let alone understand – the full range of system behavior. For these reasons, theory, experimentation, and modeling must be informed and validated by field studies. Research in complex adaptive systems must embrace the full range of scientific tools from passive observation, case studies, comparative studies, and historical studies at one extreme, to experimentation and highly abstracted models at the other.

D. Particularly in sociotechnical systems (with the exception of financial markets);

Making contact with real-world observation and experiment is nowhere as difficult but important to Sandia as in the domain of sociotechnical systems. In

this domain, research rarely embraces an adaptive perspective; instead, systems are analyzed largely within the context of their engineering drawings and their designer's intent.

One reason for this is that researchers and the public both struggle to embrace the notion that such systems, comprised of people and the technology they fashion to serve them, have adapted beyond the bounds of their blueprints and designs; acquiring independent goals and new means to achieve them. The independence and goal-seeking of an organism or group of organisms is easy to imagine, but it is not nearly so easy and natural to think about the goals of the power grid system itself (as opposed to what it was designed for, or the purposes one constituency – for example, the government – wants it to serve).

A second reason for this is that design and performance data are more precise, plentiful, and accessible than social data, and researchers are more skillful and experienced in modeling and analyzing it. Thus, even with efforts that seek to embrace the adaptive nature of sociotechnical systems, the focus often deflects toward the designed purpose and the engineered technical components of those systems and falls short of the mark.

Financial markets and economies are significant exceptions to this generalization. Markets are generally studied with the understanding that they have a life of their own. We "stimulate" markets and study how they "respond" and "behave," rather than how we will control them. Even regulators, politicians, and the media generally treat markets as adaptive systems. Research on these systems better balances social and technical aspects by presuming less complete understanding of system structure and less centralized control. Even here, however, it is rare for systems to be treated as wired, hierarchical networks of diverse and local elements adapting to achieve diverse and local goals.

The scarcity of rigorous sociotechnical research from a genuine complex adaptive systems perspective leaves us largely ignorant about the parameter space within which this perspective is most useful. For what inquiries about what systems are complex adaptive qualities significant? The usefulness of being able to answer this question extends beyond vast and rapidly changing systems like power grids, the Internet, and social media networks, to computer networks, sizable design projects, and the infrastructure of mid-sized cities. Even the design or modification of rather simple components destined to be inserted into complex sociotechnical systems might, like an organ transplant, be usefully considered as risky and unpredictable interventions into vast and mysterious self-regulating systems. Only time will tell at what aperture settings this perspective becomes more effort than it is worth.

E. Though progress in biology can be an exemplar for complex adaptive systems S&T.

Not coincidentally, biology, the field in which the most progress in complex adaptive systems research has been made, is also the field in which the most interaction and reconciliation has occurred: between theorists, experimentalists and those studying systems in the wild; and between researchers focused on functional explanations (e.g., genetics, microbiology and chemistry) and those focused on teleonomics (e.g., evolutionary biology).

Indeed, as recently as 1930, genetics and natural selection were unreconciled. Natural selection was still not accepted as the only major cause of adaptation. No one could figure out how discrete changes in genes could result in the continuous variation naturalists saw in the field. No plausible hypothesis existed to explain how genomic changes in populations resulted in adaptation and speciation.

Over the subsequent fifty years, however, the synthesis of these fields caused a creative explosion. Observation, experimentation, modeling and theory catalyzed each other to new heights. As examples, theoretical speculations by Francis Crick, Richard Dawkins and others led to genetic discoveries like meiotic drive and selfish DNA, and population models revealed critical holes in our understanding of adaptation and speciation, which led to new field observations, experiments and theories. The combinations of controlled and field studies and of functional and teleonomic explanations were critical to these advances. Today, both genetics and natural selection have so permeated all aspects of biological thinking and research that they are no longer separate areas of study. They have become part of the very fabric of ecology, developmental biology, cellular biology, animal behavior, and other domains.

Biology can thus serve as an exemplar for how to: use a domain focus to reconcile functional and teleonomic perspectives; synthesize observation, experimentation, modeling and theory; and mature complex adaptive systems research. Biological systems are orders of magnitude more complex than strictly chemical or physical systems or than any engineered system or model we could create. This makes successes in biology all the more inspiring.

5 What role could Sandia play in advancing complex adaptive systems S&T?

A. A unique and powerful one, through its access to real-world complex adaptive systems;

Complex adaptive systems S&T is ready for fuller contact with real-world observation and experiment. As an engineering laboratory steeped in engineered and sociotechnical systems and the domain-area knowledge associated with them, Sandia is well positioned to make important contributions to complex adaptive systems S&T in a number of ways.

In this section, we outline a few examples of these ways. We emphasize that these examples reflect our own biases, areas of expertise, and limitations in knowledge. They are not an attempt to be even-handed or comprehensive; for that we recommend a separate follow-on effort. But we believe they could be fruitfully used as a starting point for such an effort.

B. First, by developing the "evolutionary biology" of complex adaptive sociotechnical systems;

A first example of how Sandia could help advance complex adaptive systems S&T is to develop the equivalent of "evolutionary biology" for complex adaptive sociotechnical systems. As discussed earlier, biological and other complex adaptive systems share fundamental structural properties, and we have argued that biology can be an exemplar for S&T development in engineered and sociotechnical complex adaptive systems. However, it is not yet clear to what extent the developmental and adaptive properties of these diverse types of complex adaptive systems are similar.

Indeed, outside of biology we know very little about the evolutionary processes associated with how module modifications arise, are tested, and then artificially selected for or against. We do not know how or why layers are created and destroyed in these systems, what conditions increase and decrease variability among modules, what conditions stimulate cooperation versus competition, what affects adaptation rates, how protocols arise and gain acceptance, and what affects the alignment of the goals of elements and their sub-elements. While, for example, the proliferation of hardware and software layers in computing appears sensible in retrospect, we have little understanding of the processes by which it occurs or the conditions that encourage or discourage layer formation. And we do not know whether any of these phenomena are consistent across diverse system types. These questions have hardly been asked about sociotechnical systems – an indication of how rarely sociotechnical systems have been researched from a genuinely adaptive perspective.

Nonetheless, given the tremendous advances that have been made in evolutionary biology, we believe similar advances in "evolutionary sociotechnical systems" are well worth attempting. Indeed, some work has been done in related fields, such as technology adoption, most notably by Everett Rogers in his work on innovation diffusion.[43] And, of course, Sandia's intimate domain knowledge of how real sociotechnical systems evolve is a strength to be utilized.

Three results can be imagined.

A first result would be a new evolutionary sociotechnical framework that fuels a fertile spiral of theory and experiment. The framework would have practical implications: pointing towards potential means to stimulate or retard system growth or variance, strengthen or weaken protocols, address fragilities, enable adaptation to new environments, and stabilize or destabilize system states. It would reveal which properties are common to all complex adaptive systems and which are domain specific, enabling one to determine which techniques and principles should be borrowed from other domains.

A second result would be an ability to discern when adaptive dynamics can or cannot be safely ignored, and hence how narrow or wide an aperture one needs to take in considering a given complex adaptive system. Sociotechnical systems are subject to unique outside interventions, such as government regulation, tax incentives, wholesale redesign, and shutdown and restart. The effects of these interventions on adaptive behavior should also be understood. A beginning could be made by exploring the conditions affecting the rates of different types of adaptation and the effects of artificial selection, regulation and other anthropogenic factors on the system.

A third result would be an ability to predict, within bounds, the rate and path of adaptation in various kinds of environments. For conventional (non-complex-adaptive) systems, such uncertainty quantification (UQ) and risk assessment (RA) is a key strength of Sandia. In these systems, the probabilistic interpretation of confidence in, say, an intervention that prevents a real-world complex adaptive system from reaching a security threshold, is formally understood. For complex adaptive systems, however, UQ and RA cannot be performed as if engineering

---

[43] E.M. Rogers, "Diffusion of Innovations" (Free Press, 1st Ed, 1962, 5th Ed, 2003).

systems are isolated entities, but must recognize the human response to the engineering system's use or failure. In this context, the advance of UQ and RA methods for complex adaptive systems would be an important extension to its conventional application to engineering systems.

C. Second, by developing new abstraction and analysis techniques for complex adaptive systems;

A second example of how Sandia could help advance complex adaptive systems S&T is in developing new techniques for abstracting and analyzing complex adaptive systems.

As discussed in section 4, such abstraction and analysis is difficult because the reductionist approaches associated with the physical sciences has limited applicability to complex adaptive systems. On the one hand, detailed modeling is powerful, but does not provide clarifying simplicity and does not scale well in systems spanning many space and time scales and composed of many layers and modules. On the other hand, mean-field approaches are also powerful, but abstract away structures essential to system function. Neither approach focuses on the adaptive signals available in these systems.

The challenge, as with all modeling, is to capture the right qualities, relationships and functionality with sufficient accuracy for the question at hand and to leave out the rest. Models must simplify in order to be useful, and it seems likely that any major mathematical breakthroughs in complex adaptive systems will involve some hybrid of detailed modeling and mean-field techniques with newer approaches designed to leverage adaptation.

Here, we suggest three analytic approaches, based on leveraging known adaptive features such as those described in section 3: the mutual fit of structure, function and environment due to past adaptation; the coevolution of adaptive elements; and the simplicity, ubiquity, stability, and exploitability of protocols.

A first analytic approach is to leverage the mutual adaptation of structure, function, and environment. Knowledge of mechanisms may be used to understand higher functions and vice versa. This is, in part, a typical reductionist approach from the physical sciences, leverage phenomena at one scale to understand it at another (e.g., understanding what molecular structures result in higher viscosities).

But, in complex adaptive systems, it is also more than reductionism. In the physical system case, there was no demand for viscosity. It was not survival of the fittest, but merely survival of the stable. In contrast, adaptive structures were selected for their functionality. Thus knowledge of a system's function allows one to hypothesize means and vice versa.[44] For example, knowing what the holistic purpose of the kidney as a system embedded in its human body environment gives clues to what reductionist biochemical pathways are important for the functioning of the kidney in accomplishing its purpose.

Note that mechanisms need not be below a function they serve or even in an adjacent layer. In our bodies muscle activity causes chemical changes at lower levels, which we experience as soreness, which in turn influence our level of physical activity. In a city, transportation enables growth which demands more transportation. Multilayer interactions are a hallmark of hierarchically modular structures, and complex systems engineers must be equally adept at working downward (reductionism) and working upward (holism). The physical sciences' assumption that the basement layers of a system are more fundamental and causal than higher layers is unjustified in complex adaptive systems.

One of the challenges with this analytic approach is that the mathematics of teleonomy and adaptation are immature compared to those of functional causation. We can follow a single flow of energy or information along a functional chain and combine it with the effects of other flows, even, to a large extent when these flows feed forwards and backwards into themselves and each other. Following the multiple threads associated with mutual adaptedness is more difficult, though one can anticipate borrowing techniques from genetics, ecology and population biology, in addition to using control theory, machine learning, agent-based modeling, and other techniques.

If successful, advances in this arena will enable improved methods for addressing "why" questions, such as why inner cities collapse, why social groups form and dissolve, and why data structures and software evolve as they do; they will facilitate rapid discovery of functions and mechanisms in unexplored or even obscured or largely inaccessible systems; and they will enable modelers to better distinguish critical structures and dynamics.

---

[44] An example of how system functionality can be quantified and exploited mathematically is John Doyle's Robust-Yet-Fragile (RYF) theory and Sandia's practical advances upon it. See, e.g.: J.M. Carlson, J. Doyle, "Complexity and robustness," Proc. Natl. Acad. Sci. USA **99**, 2538-2545 (2002); and R.A. LaViolette, K. Glass, R. Colbaugh, "Deep information from limited observation of robust yet fragile systems," Physica **A 388**, 3283-3287 (2009).

Further, it is possible that, by focusing first on function rather than structural detail,[45] broadly applicable structural motifs that give rise to function will be discovered. In biology, certain solutions so elegantly and robustly solve a recurring functional problem that they are discovered time and again by natural selection, from protein synthesis processes at lower levels to larger scale structures such as eyes or social pecking orders. Similarly in engineering, different cultures discovered arches and triangles for structural engineering; and wheels, gears, and levers for mechanics. There may be many such motifs in adaptive sociotechnical systems yet to be discovered—network structures, logical structures, dynamical patterns. Like simple machines, these motifs will obey physical laws and make clever use of them, but will not themselves be physical laws.

A second analytic approach is to leverage co-evolution and other signals to anticipate future adaptation.

Adaptation, of course, is highly constrained. Frequently represented on a multi-dimensional and changing fitness landscape, adaptation normally occurs by incremental steps. Although the design space is vast, the possible incremental changes to the structure and environment that leave the system functional – let alone improve performance – are generally a minute portion of the combinatorial space. (In biology, mutations are almost always deleterious.) The viable options are highly dependent on legacy design choices, the current environment, and available internal variation. It is easy to imagine beneficial adaptations – arms for birds, for example – that are not reachable designs, as these systems lack much foresight and are highly biased against a descent into a valley that could lead to escape from a local maximum. Thus current and past functional design both drives and constrains future adaptation.

Although we know it *is* constrained, we don't yet have the techniques to quantify *how* it is constrained. One possibility is to make use of "teleonomic" techniques, which are proving useful as augmentations to existing signal processing and modeling. Processing a co-evolving signal as if it were any other data stream is like teaching someone a game without mentioning the goal of the game or the motives behind any actions.[46] Another possibility is to make use of signals associated with individual variation, e.g., among modules. Teleonomic processes often produce populations of similar entities with important differences or variances.[47] Variance and circumstance lead to specialization, which is the seed for speciation, a permanent division of the group. This specialization matters in day-to-day performance and provides important signals about the potential for and direction of longer-term adaptation.

Finally, a third analytic approach is to make use of protocols to determine the macroscopic characteristics of complex systems inexpensively. Protocols may govern how nodes and links are formed and be the control system for of traffic among them. Thus protocols provide shortcuts to understanding topology and dynamics.

For all three analytic approaches, it will be important to build on approaches developed in other domains. Reachability, control theory, genetic algorithms, and other optimization techniques offer promising analytics. Techniques from biology, including population genetics and evolutionary game theory should prove useful. New knowledge about the adaptive processes themselves (see section 5B) may suggest other approaches. And techniques associated more recently with complex adaptive systems (control theory, reachability analyses, game theory, genetics, population biology, machine learning, dynamical systems (Lyapunov, Markov), agent-based modeling) should also be useful.

D. Third, by leveraging social theory and digital social data to advance our understanding of human behavior;

A third example of how Sandia could help advance complex adaptive systems S&T is to better understand

---

[45] In one example, Doyne Farmer at the Santa Fe Institute is having success modeling financial markets while assuming "zero-intelligence investors," i.e., that investor strategies, in the aggregate, are random [J. Doyne Farmer, P. Patelli, I. Zovko, "The Predictive Power of Zero Intelligence in Financial Markets," Proceedings of the National Academy of Sciences, (February 8, 2005)]. In another example, Rich Colbaugh at Sandia predicts total movie box office while assuming that the quality of the movies is unimportant. Farmer's model details order flows and Colbaugh's model details the social buzz about the movie. Each combines detailed transactional model elements (reductionism) with averaging approaches. Each also emphasizes behavior at middle scales. In Colbaugh's case important meso-level phenomena are the buzz passing between social communities and whether the buzz passes through central graph clusters.

[46] Recent work by Rich Colbaugh shows the power of accounting for these adversarial goals in spam filtering.

[47] For example, a market consists of buyers and sellers who tend to occupy unique niches as they pursue their commercial goals. Similarly, a biological species or local population is a diverse collection of individuals, adapting individually and collectively at multiple levels and timescales.

one of the most important aspects of complex adaptive systems: the presence of humans.

With humans, the central challenge is to move beyond the "humans as exogenous to the system" approach[48] preferred at Sandia and elsewhere up until recently, towards a "humans as endogenous to the system" approach. The principal advantage of the humans-as-exogenous-to-the-system approach is that system problems are circumscribed to the purely engineering domain that is Sandia's expertise, and remain manageable. Its principal disadvantage is that the engineered system, though it might meet its narrow engineering objectives, risks failing to achieve its ultimate social objective (e.g., deterrence). In a worst case, the engineered system might end up being socially unusable or unacceptable, or might be compromisable through social means not anticipated or addressed during design.

Our belief is that realistic treatments of a wide range of systems *must* include humans as endogenous to the system, i.e., must include people as essential adaptive components, as it is ultimately people whose needs and wants are being served.[49] To keep the human exogenously in the environment forces us to miss the complex interplays of human behavior modifying the system and the system modifying human behavior, and to miss the emergent behavior of the system as a result of those modifications.

Moreover, we see the glimmers of quantitatively being *able* to treat humans as endogenous to the

system. Now, for the first time, it may be possible to begin treating systems containing humans like they truly are: complex. By enabling the collection and analysis of fine-grained social data from people "in the wild,"[50] the digital revolution is making possible quantitative validation of social theory. Recent advances in deformable electronics[51] may even enable humans to be outfitted with a new generation of intelligent, communicative sensors, thus enabling a new level of detail and ease in the collection of data.

In addition, human behaviors can be computationally modeled based on well-vetted psychological-social-behavioral and economic theories[52] which account for the uncertainties of the model inputs, the human model, and the model

---

[48] Note that this approach has, within its limitations, worked quite successfully in the past (as in the classic "man machine method material" interacting components model made famous by Kaoru Ishikawa in his root-cause diagrams, see, e.g., Kaoru Ishikawa, "Introduction to Quality Control" (Productivity Press, 1990)). By treating human beings as exogenous to the system (i.e., as part of the environment), human and social "factors" are regarded as an interface between the engineered system and the human and social environment the system is embedded in. The requirements of human users are solicited and treated as exogenous inputs to and outputs of the engineered system (see, e.g., J.P. Brannen and K.L. Hiebert-Dodd, "Expanding the domain of systems analysis," Simulation **49**, 141 (1987)).

[49] Moreover, this is not just our belief. The Department of Defense (DoD) has, in the last ten years, dramatically increased its emphasis on hearts and minds at the expense of "shock and awe." National security is increasingly being viewed as a people issue, with hardware being just the tools that people use. Indeed, as DoD continues its shift away from the Fulda Gap and toward stability operations, we believe it inevitable that the Departments of Homeland Security and Energy will begin their own similar shifts. Ultimately, it is the humans-as-endogenous-to-the-system approach that will dominate.

[50] See, e.g., the Computational Social Science Initiative based at the University of Massachussets Amherst (http://cssi.umass.edu/index.html), and David Jensen's keynote lecture at the KDD-2010 (Knowledge Discovery and Data Mining) conference (http://videolectures.net/kdd2010_jensen_css/).

[51] D.H. Kim, N. Lu, R. Ma, Y.S. Kim, R.H, Kim, S. Wang, J. Wu, S.M. Won, H. Tao, A. Islam, K.J. Yu, T.I. Kim, R. Chowdhury, M. Ying, L. Xu, M. Li, H.J. Chung, H. Keum, M. McCormick, P. Liu, Y.W. Zhang, F.G. Omenetto, Y.G. Huang, T. Coleman, J.A. Rogers, "Epidermal Electronics," Science **333**, 838-843 (2011). See also Z.Q. Ma, "An Electronic Second Skin," Science **333**, 830-831 (2011).

[52] I. Ajzen, "Theory of Planned Behavior," Organizational Behavior and Human Decision Processes Volume **50**, Issue 2, December 1991, Pages 179-211; I. Ajzen, M. Fishbein, "Attitudes and the Attitude-Behavior Relation: Reasoned and Automatic Processes" European Review of Social Psychology (2000) Volume: **11**, Issue: 1, Pages: 1-33 (history and update of the Fishbein Azjen work on their expectancy-value model of attitude); R.E. Petty and J.T. Cacioppo, "The Elaboration Likelihood Model of Persuasion," Advances in Experimental Social Psychology (1986) Volume **19**, Pages, 123-187; L. Festinger, A Theory of Cognitive Dissonance (Stanford University Press, 1957); A. Bandura, "Social Learning Theory of Aggression," Journal of Communication,Volume **28**, Issue 3, pages 12–29, (1978); H.A. Simon, "Bounded Rationality," Organizational Science, Volume **2**, Number 1, pages 125-134 (1990); D. McFadden, "Conditional Logit Analysis of Qualitative Choice Behavior," Chapter 4 of Frontiers in Economics, Academic Press, Pages 105-142 (1973); J.E. Stiglitz, "Equilibrium in Product Markets with Imperfect Information," The American Economic Review, Vol. **69**, No. 2, Papers and Proceedings of the Ninety-First Annual Meeting of the American Economic Association, pages 339-345 (1979); D. Kahneman, J. L. Knetsch and R. H. Thaler, "Experimental Tests of the Endowment Effect and the Coase Theorem," Journal of Political Economy, Vol. **98**, No. 6, pages 1325-1348, (1990); C.W.J. Granger, "Causality, cointegration, and control," Journal of Economic Dynamics and Control, Volume **12**, Issues 2-3, Pages 551-559, (1988).

outputs (behaviors). What previously were untestable hypotheses can now be subjected to the scientific method and occasionally elevated to the status of verified knowledge.

Finally, social theory is maturing; neuroscience is advancing our understanding of the brain; techniques have been developed to explain and predict cognition and behavior in certain circumstances; and the digital age has provided goldmine of social data. Sandia has built computational models based on human cognition[53], improved and extended the use of agent-based models, and exploited meso-scale network structures to predict group behavior.

Taken together, we see dramatic advances in computational social science, an age that Sandia could both benefit from and contribute to. This will have deep ramifications on the design and engineering of sociotechnical systems.

That said, a national laboratory like Sandia must approach human complexity with caution. We are entrusted with a duty to serve the nation, not to manipulate it through the design and engineering of systems whose goals have a social engineering or influence component to them. We believe these are exceedingly difficult issues, ones which must be deliberated, debated, and continually revisited as Sandia threads the needle between performing its service to the nation and avoiding even the appearance of abuses en route to providing that service.

Sandia has proactively initiated principles and guidelines[54] for developers and we list three minimal actions that will be essential in our LDRD work.

First, we should clearly and openly articulate why our service to the nation requires expanding our systems view to include humans as integral components. The reasons should include both the general reasons given above as well as concrete clarifying examples. Deterrence is a clear example of a critical mission whose ultimate goal is social influence. Second, we should initiate a broad-ranging debate to define the "line in the sand" that demarcates systems that we will not design or engineer because they could be too easily co-opted for nefarious social manipulation.

Third, since people both make and are the object of policy, we need to articulate clearly the line between policy (and decision) support, which Sandia does, versus policy advocacy, which Sandia does *not* do.

E. Fourth, by developing new observational and experimental platforms and techniques for complex adaptive systems.

A fourth example of how Sandia could help advance complex adaptive systems S&T is in the development of new observational and experimental platforms and techniques for complex adaptive systems. As indicated in figure 5, both natural (field studies) and artificial (experiments) are important platforms, but they fill very different roles.

Field studies and natural experiments, while the least controlled, reveal the natural behavior of real complex adaptive systems. These are common tools for biologists and social scientists, and are becoming increasingly common tools for complex adaptive systems scientists. They can also span a wide range of domains: from ant colonies and bee hives that can be maintained in highly controlled environments; to mobile phone networks that free-run in uncontrolled environments.

Artificial experiments, while the most difficult, may hold the most long-term promise, for two reasons.

First, they can be made simple and designed to test unambiguously particular theories and hypotheses: like controlled experiments applied to physical systems characterized by disorganized complexity, but applied to complex adaptive systems with organized complexity. For example, in the 1980's Robert Axelrod ran a computer program Prisoner's Dilemma tournament and gained key insights into the evolution of cooperation among animals, individuals and organizations. More sophisticated experiments for exploring sociotechnical systems are being enabled by emerging social media technologies.

Second, they can take advantage of the intentional modification of key components, to test how those modifications influence the behavior of the system. Indeed, one can take advantage of a key transformation that modern civilization is undergoing: from a "components are stupid" paradigm to a "components are engineered sensored, intelligent and communicative modules" paradigm – a paradigm in which components can rewire both themselves as well as their interconnections with other components in response to their recent functional interactions with their environment.[55]

[53]M.L. Bernard, M. Glickman, D. Hart, P. Xavier, S. Verzi, P. Wolfenberger, "Simulating Human Behavior for National Security Human Interactions," SAND2006-7812, Albuquerque, NM: Sandia National Laboratories (2006).

[54]W. L. Shaneyfelt, "Ethical Principles and Guidelines for the Development of Cognitive Systems," SAND2006-0608, Albuquerque, NM: Sandia National Laboratories (2006).

[55] These smart machines have obvious parallels and synergies with humans as modules in complex adaptive systems. On the one hand, machines are not humans, so

For both field studies and artificial experiments, data will be key – particularly the massive amounts of data[56] associated with massive numbers of modules and sensors or actions associated with the modules. Each adaptive element in a system collects and retains data, sends and receives information or material, processes these internally, and modifies its own behavior and structure. In the aggregate, these changes cause the evolution of the system as a whole. Even fine-grained transactional data from complex systems captures only a fraction of these phenomena.

This is all the more the case when the time dimension, which is critical for adaptive systems, is considered. These systems exhibit dynamic behaviors and they modify their own structure over time. Heterogeneous behaviors and modifications occur at multiple system levels. The useful metadata, or global measures, of such systems have yet to be developed. For example, existing measures of networks, including diameter, average degree, centrality, and betweenness, are largely static. Better means of capturing the state of a dynamic network and its modification over time are needed, for example measurements of the strength of dependencies or relationships. Information theory, which addresses only the traffic of bits, should be extended to address their effect on the states of senders and receivers, i.e., communication and its effects.

Thus, a major challenge will be to understand what types of data to collect, and what not to collect, enroute to meaningful real-world or controlled experiments.

Finally, we re-emphasize that Sandia's strength is its access to many real domain-area complex adaptive systems. Data from these systems can be cached or fed live into experimental environments, or these systems can be co-opted temporarily for research purposes. However, this is counter to the normal engineering and problem-solving mind-set at Sandia. It is in the nature of an engineering laboratory to prefer to manipulate the system to produce a desired functional outcome for a particular application rather than for the express purpose of testing and refining a scientific hypothesis.[57] Nevertheless, we believe this is precisely where Sandia can make powerful contributions.

6   How might Sandia invest in complex adaptive systems S&T?

Currently, Sandia has relatively little complex systems S&T capability that it developed on its own, and hence has a relatively low global profile. Moreover, it has almost no programs intended to enable Sandia to contribute to complex adaptive systems S&T. Bootstrapping a virtually non-existent activity into a world-class capability will take time and patience.

Here, we propose a three-pronged approach to building such a capability. Capabilities are built through programs in which real work is done. The trick is to find customers for those programs. This is particularly challenging in areas, such as in this one, in which Sandia is boot-strapping its way up. The three prongs are illustrated in Figure 5, and we discuss them in turn.

A.   First, as a fast follower, applying existing S&T to real domain-area complex adaptive systems;

A first prong is to continue our work as fast followers, applying advances worldwide in systems S&T to system improvement in particular domain areas. After all, the ultimate goal of complex adaptive systems S&T is to make these systems more useful to humanity. Our domain-area customers would benefit from more than just understanding of complex

---

their behaviors will be different from those of humans. Because, unlike humans, they are programmable, they offer the ability to tailor that programming to elicit particular system behaviors. In most situations those behaviors will be those considered desirable to the ultimate functioning of the system. However, it should also be possible to elicit behaviors intended to test hypotheses about the relationship between module and system behavior. Moreover, since privacy is not an issue with machines, observational and experimental data can in principle be much more direct and detailed. On the other hand, machines can be made to emulate particular aspects of humans. Thus, they can be used to test particular hypotheses associated with human behavior and how that behavior influences overall complex adaptive system behavior. They can also be used in a biomimetics manner: using our understanding of humans to improve the design of intelligent machine modules. Sandia's combined expertise in robotics and in human cognitive science position Sandia well to bridge between these two communities and to explore similarities and differences in how smart module microbehavior rolls up into overall system macrobehavior.

[56] A. Chervenak, I. Foster, C. Kesselman, C. Salisbury, S. Tuecke, "The data grid: towards an architecture for the distributed management and analysis of large scientific datasets," J. Network and Computer Applications **23**, 187-200 (2000).

adaptive systems; they would benefit from an ability to build, influence, leverage, and improve them.

Currently, these domain-area customers are dominated by the Department of Defense and the Department of Homeland Security, but it will be important to begin socializing the possibility of similar programs supported by the technology offices in the Department of Energy.   Though the Department of Energy has not recently supported S&T in complex adaptive systems, the time may be ripe for it to move in this direction.

Note, though, that there are three very distinct routes to system improvement in a particular domain area. All three routes can be important to Sandia, to its Department of Defense and Department of Homeland Security customers, and to its potential Department of Energy technology office customers.

The first two routes are system improvements that occur over an engineering design and implementation, or "evolutionary" time scale.

The first of these "evolutionary-time" routes to system improvement is direct: to improve the system as a whole through a rewiring of the interactions between existing sub-systems and components, or through very minor "on-off" modifications of these sub-systems and components.   This route depends on a deeper understanding of how the sub-systems and components interact and conspire to give the system its overall performance, and is the current center-of-gravity for Sandia's systems and complex adaptive systems work.[58],[59],[60],[61],[62]   For example, one might rewire a social contact network through forced shut down of particular hubs (schools, churches) or portals (public transportation) in order to contain a pandemic in a complex adaptive social network.[63]

The second of these "evolutionary-time" routes to system improvement is indirect: by enhancing functionality and performance of the components themselves. Note that this route normally depends solely on a deeper understanding of component performance, because the component normally interacts with the rest of the system through a rigid protocol.[64]   However, the potential system improvement could be expanded if the protocol itself could be intelligently shifted through a deeper understanding of the larger system.[65] In other words, we believe there is considerable opportunity for "what if" module performances based on "what if" system designs. If a module researcher knew that improving module performance in a certain manner would enable boundaries between modules to be shifted with a resulting improvement in overall system performance, more effective investments into module research could be made.   There is opportunity for system design and engineering to inform, in a much deeper way than it has in the past, investments into subsystem and component research.

The third route is system improvement that occurs through real-time influencing of system status, goals and outputs. That is, to move beyond "evolutionary-time" design improvement to "real-time" nudging (and perhaps even command and control): towards stability or instability, towards cooperation or competition, towards convergence or divergence of goals, or even towards new protocols.   Such nudging will require instrumentation for near-real-time monitoring, analysis, visualization, and decision support.   It will require recognizing stable and unstable states of these systems and desirable and undesirable "nearby" states. It will require predictive

[58] R.J. Glass, W.E. Beyeler, S.H. Conrad, N.S. Brodsky, P.G. Kaplan, T.J. Brown, "Defining Research and Development Directions for Modeling and Simulation of Complex, Interdependent Adaptive Infrastructures," SAND 2003-1778, Albuquerque, NM: Sandia National Laboratories (2003).

[59] R. LaViolette, W.E. Beyeler, R.J. Glass, K.L. Stamber, H. Link, "Sensitivity of the resilience of congested random networks to rolloff and offset in truncated power-law degree distributions," Physica A:Statistical Mechanics and its Applications, Vol 368, Issue 1, pp 287-293 (August 2006) (see also SAND2005-5926 J).

[60] W.E. Beyeler, R.J. Glass, M.L. Bech and K. Soramäki, "Congestion and cascades in payment system," Physica A, Volume 384, Issue 2, pp 693-718, (2007).

[61] R.J. Glass, T.J. Brown, A.L. Ames, J.M. Linebarger, W.E. Beyeler and S.L. Maffitt, "Phoenix: Complex Adaptive System of Systems (CASoS) Engineering Version 1.0," SAND 2011-3446, Albuquerque, NM: Sandia National Laboratories (2011).

[62] S.H. Conrad, W.E. Beyeler, T.J. Brown. "The Value of Using Stochastic Mapping for Understanding Risks and Tracing Contaminant Pathways. Proceedings of the 4th Annual Conference on Infrastructure Systems: Challenges and Research for Next Generation Infrastructures in the 21st Century" (in press) (2011) (see also SAND2011-4203C).

[63] V.J. Davey, R.J. Glass, H.J. Min, W.E. Beyeler, L.M. Glass, "Effective, Robust Design of Community Mitigation for Pandemic Influenza: A Systematic Examination of Proposed U.S. Guidance," PLoS One, 3(7): e2606, doi (2008).

[64] For example, one might improve the efficiency of a white-light illumination source but otherwise maintain its output color temperature at the usual 2700K.

[65] For example, if one understood how illumination sources with different color temperatures affect human mood, Circadian rhythm, perception of brightness and visual acuity, one might develop illumination sources having alternative color temperatures and thereby improve overall human productivity.
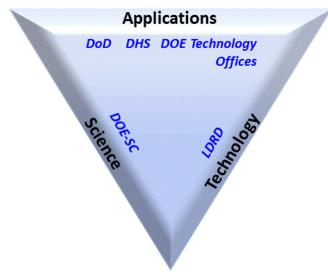
Figure 5. Interacting triangle of science, technology and applications, along with current or potential customers for programs in those areas. DoD = Department of Defense; DHS = Department of Homeland Security; DOE = Department of Energy; SC = Office of Science; LDRD = Laboratory Directed Research and Development.

analytics and scenario (what if?) analytics. And it will require evaluation of how early, accurate, and specific a given signal is under varying circumstances. The tradeoffs of more and fewer sensors and more or less frequent data capture must be evaluated. As humans will continue to make most decisions, visualization, uncertainty quantification, and the effects of these on human decisions must be studied. The ultimate engineering goal would be to learn to "grow" a complex adaptive system to perform an adaptive function in a variable environment.

We believe all of these routes can be important. The first route we do already. The second and third routes we don't yet do, but if we did could expand our sphere of impact significantly.

B. Second, contributing to S&T itself through internal LDRD programs;

A second prong is to contribute to complex adaptive systems S&T, as opposed to simply being world-class "appliers" of existing S&T to complex adaptive systems domain-area problems. The reason: contributing to the S&T provides the marginal *knowledge*, *people*, and *differentiation* necessary to credibly tackle the problems that Sandia should be tackling, those that are most difficult and challenging for our nation. It is these "national laboratory" class problems that if we don't tackle won't get tackled.

With regards to *knowledge*, contributing to S&T provides two distinct benefits. First, it provides us with our own foundational S&T base, a base for

translational[66] S&T that can be used to solve real domain-area problems. Second, by playing "on the same field" as other contributors to foundational S&T, we gain better access to their cutting-edge work.[67]

With regards to *people*, contributing to S&T is a powerful recruiting tool. Unlike problem solving, S&T is almost always published, and if it is of a high quality, serves to attract outstanding scientists and technologists. These scientists and technologists would in turn strengthen our S&T base, and some would go on to apply that base to real problems.

With regards to *differentiation*,[68] contributing to S&T would provide Sandia and its energy and national security customers with complementary and mutually reinforcing perspectives. On the one hand, Sandia's perspective could become one in which it preferentially tackles the most difficult problems of national interest – those problems which require a deep S&T base – leaving other problems to other laboratories to tackle. In other words, it would provide a criterion for Sandia to choose which problems to tackle. On the other hand, Sandia's national security customers' perspective could become one in which they view Sandia as being uniquely capable of tackling their most difficult problems of national interest. These customers would come to differentiate our capabilities from those of the other laboratories, based on a difficult to attain core S&T strength.[69]

We have to recognize, though, that this is difficult to pull off, even in areas that rest on a relatively narrow

---

[66] Indeed, we advocate, as we discussed in section 3, bi-translational S&T: making use of S&T to solve domain-area problems, but also using domain-area problems as vehicles to contribute new S&T.

[67] Only by participating in cutting-edge S&T oneself can one self-reinforce the absorptive capacity and/or the credibility necessary to participate in the formal and informal network of cutting edge S&T.

[68] In the resource-based view in business management [B. Wernerfelt, "A resource-based view of the firm," Strategic Management Journal **5**, 171-180 (1984)], one can think of this as building a sustained (as opposed to short term) competitive advantage that is valuable, rare, inimitable (or imperfectly imitable), and non-substitutable. Branding and communications can help with the appearance of meeting each of the above criteria, but true sustainable competitive advantage can't be based on messaging alone, but rather on real capability.

[69] Note that an alternative path towards differentiation might be as fast followers in S&T, through our access to and understanding of domain-area problems and of the multi-disciplinary S&T necessary to address those problems. Without significant barriers to entry, however, this path may not provide sustainable differentiation.

knowledge base. Moreover, even when achieved, the ability to integrate cutting-edge science with cutting-edge engineering is not trivial: cultural divides between the scientific and engineering communities are common and difficult to surmount. All that said, the rewards are significant the problem space that is opened up is wide and deep, and the impact that Sandia could have on the nation is exceptional.

In proposing that Sandia contribute to S&T through internal programs, and that Sandia boot-strap an S&T capability nearly from scratch, we note that there are two possible paths, each with its own strengths and weaknesses.

An extreme bottom-up approach, determined by individual initiative and entrepreneurship, is the usual default in a brand-new area too small to catch management attention. An extreme top-down approach, determined by strategic plan, is the usual default in an established area so large that management cannot afford investment mistakes.

Their strengths and weakness are complementary. The strength of the bottom-up approach is the unleashing of the ingenuity of individual staff and their in-the-trenches appreciation of where the best opportunities lie to make research contributions; the weakness of the top-down approach is the absence of this unleashing. The weakness of the bottom-up approach is the absence of a "greater-than-the-sum-of-its-parts" cross-organizational synergy that is one of Sandia's great strengths; the strength of the top-down approach is the thoughtful building of this synergy.

Here, we propose a combination of the two approaches, each modified to allow for some elements of the other.

Our proposed bottom-up approach recognizes that Sandia is already working on a wide set of complex adaptive systems problems in specific domains areas (national security, homeland security, energy, climate and infrastructure security), problems that could easily be the springboard for developing scientific insights and technological tools that would be generalizable across domain areas. We propose to give our research staff the freedom, perhaps even the mandate, to develop such generalizable scientific insights and technological tools, without specifying too narrowly the nature of those insights and tools. This kind of bottom-up selection would naturally lead to the best insights and tools, the kind that cannot be anticipated ahead of time.

Thus, we propose a bottom-up approach based on small internal LDRDs aimed at contributing to S&T capability that either: (a) could make a long-term impact on two or more of these domain-specific programs; or (b) uses one or more of these domain-specific programs as a vehicle for research. Here, we have in mind a bottom-up call for proposals, in which the S&T area is not specified, except in relation to its bi-translational synergy with our existing domain-specific complex adaptive systems programs.

Note that even though we cannot specify the particular insights and tools that might actually be developed, we can say something about the class of insights and tools that would maximize benefit to Sandia and the nation. It is precisely those insights and tools that are generalizable across the domain areas already known to be of interest to Sandia. This is not too restrictive, as Sandia is already interested in a wide range of domain areas. But it is restrictive enough, as the very definition of bi-translational S&T is that it have broad applicability across domain areas.

Thus, we propose to augment the bottom-up approach with a top-down approach that recognizes that some S&T areas are so obviously important, synergistic across organizations, and differentiating that they can be pre-identified strategically. Indeed, if they were not so pre-identified, one might expect them to be under-invested in through a simple bottom-up approach. This top-down approach would be a laboratory-wide grand challenge LDRD aimed at building S&T capabilities in the areas identified in section 5. Here, we have in mind a large program structured around the top-down areas of S&T identified in section 5, and consisting of cross-organizational teams. This Grand Challenge, along with the smaller LDRDs discussed above, would help build Sandia's reputation for complex adaptive systems analysis research. An External Advisory Board would be critical for steering and recalibrating the actual S&T areas, and for facilitating customer connections and enable interactions with academia.

C. Third, by developing long-term support for complex adaptive systems science through the DOE Office of Science.

Finally, we note that even in the physical sciences, Sandia relies heavily on outside sources (particularly the DOE Office of Science) for support for its contributions to science. Thus, we believe it will be beneficial to begin socializing the possibility of complex adaptive systems science to DOE-SC, perhaps with a Basic Research Needs for Complex Energy Systems Science Workshop. This will be key to Sandia's success in this area; without separate long-term funding for complex adaptive systems science, it will be very difficult for the current set of applications-driven customers to support science and technology capability development. This is perhaps also the most difficult prong. DOE-SC (the Office of Science)

currently focuses on physical sciences and does not have a complex systems or system of systems research focus. On the other hand, the successful market acceptance of new energy technologies depends on the complexity rich environment where they will be implemented and on the societal functions the technologies serve. DOE may be receptive to this incremental perspective on addressing the chasm between technological advances and commercial realization (the "Valley of Death").

One possibility is, if we are successful at developing programs for the DOE technology offices (EERE, OE, etc.), to use these as leverage to encourage the Office of Science to fund complex adaptive systems science. Our DOE customers may be able to make the argument on Sandia's behalf, if it can clearly be shown that such research would improve our ability to solve problems for the DOE.

Another possibility would be to illustrate that complex adaptive systems science is rapidly becoming a widely accepted and respected field. The study of complex engineered systems is going "mainstream." For example, the International Council on Systems Engineering (INCOSE), one of the principal professional societies for systems engineering, has a Complex Systems Working Group aimed at improving how systems engineering is done by applying advances in chaos, complexity, complex adaptive systems, nonlinear static and dynamics, networks, social science, neuroscience, evolution, power laws, ecology, and other topics.

D. Finally, through establishing or joining a complex adaptive systems S&T community.

In addition to following this path toward building sustainable research programs in complex adaptive systems analysis, we suggest building internal and external communities through project collaboration, technical working groups, well-advertized lectures and discussion-groups, and hosting and attending conferences. University partnerships will be key to ensuring that Sandia stays current with cutting-edge research and participation in relevant consortiums. An internal "community of practice" site could facilitate teaming on funding opportunities and the sharing of relevant research results, along with serving as a long term mechanism to identify research needs and capabilities common to many organizations and problems.

Additionally, we may need accelerate our reputation for research in this area through publication, collaboration and strategic hiring. Adding prominent researchers to the Sandia staff could help build the capability internally as well as drawing specialists into the lab and building our connections to

influential organizations, groups and publications. The concept of forming a broad network of specialists to encourage development and improvement of the discipline of complex adaptive system for addressing national security problems is addressed in detail in the Phoenix Initiative.[70]

7 Glossary

This glossary of terms is intended to help the larger Sandia community, steeped in physical science terminology, understand the nuances of the emerging language of complex adaptive systems.

Coevolution. Coevolution is the mutual adaptation of two or more elements. In other words it is a chain of adaptive change among elements over time. It is primarily a biological concept involving two or more species that reciprocally impact each other's evolution[71], but can be extended to other fields as we have in this paper.

Complex adaptive system. A complex adaptive system is a system whose structure modifies to enable success in its environment. To accommodate such adaptation, its structural organization is almost always hierarchically modular.

Emergence. The appearance of a phenomenon at a aggregate level that cannot be described in terms of the particulate actions. This is a concept used to describe non-reducible properties of systems such as community-scale patterns, whole-system function, markets, social norms.

Game theory. The study of mathematical models of conflict and cooperation between intelligent, rational decision makers.[72]

Highly optimized tolerance. A concept introduced in 1999 by Carlson and Doyle[73] to explain a mechanism for generating power law distributions in biological and engineered systems. "..systems which are optimized, either through natural selection or engineering design, to provide robust performance despite uncertain environments. We suggest that

---

[70] R.J. Glass, T.J. Brown, Ames AL, J.M. Linebarger, W.E. Beyeler and Maffitt SL, "Phoenix: Complex Adaptive System of Systems (CASoS) Engineering Version 1.0," SAND 2011-3446, Albuquerque, NM: Sandia National Laboratories (2011).

[71] http://evolution.berkeley.edu/evosite/evo101/IIIFCoevolution.shtml.

[72] R. Myerson, "An Introduction to Game Theory," Northwestern University (1984).

[73] Carlson, and Doyle. "Highly optimized tolerance: A mechanism for power laws in designed systems," Physical Review E**60**(2): 1412-1427; (August 1999).

power laws in these systems are due to tradeoffs between yield, cost of resources, and tolerance to risks. These tradeoffs lead to highly optimized designs that allow for occasional large events."

Level or layer. One row in a hierarchy. Level differences often involve different scales or a translation between incommensurate units, e.g., logic and code lines in software to electrical and magnetic states in hardware, or both.

Mean-field. A method used to analyze a system with many interacting entities whereby all the entities are replaced by one aggregate entity with the averaged properties.

Module. A portion of a system layer whose components are highly connected to each other and more loosely connected to elements in the rest of the system, i.e., a cluster. Module boundaries may be formal and governed by input/output parameters and protocols or less discrete.

Modularity. The ability for a system's components to be substituted and/or recombined, with components being able to operate independently.

National Infrastructure Simulation and Analysis Center (NISAC). The National Infrastructure Simulation and Analysis Center[74] is a modeling, simulation, and analysis program within the Department of Homeland Security (DHS) comprising personnel in the Washington, D.C., area, as well as from Sandia National Laboratories (SNL) and Los Alamos National Laboratory (LANL).

Organized and disorganized complexity. As discussed in the text, functional systems which adapt take on a hierarchically modular structure. This kind of structure is a type of "organized complexity,"[75] in which the interactions between the parts are non-random, or correlated. In contrast, non-functional systems such as a molecular gas, a sandpile,[76] or a planet's climate, are characterized by "disorganized complexity," in which the interactions between the parts are random and non-correlated. Interestingly, organized complexity is largely the domain of engineers and biologists, while disorganized complexity is largely the domain of physicists.[77]

Reductionism and holism. When we talk about the properties of a system, we are almost always talking about "holistic," macroscopic properties of the system as a whole in relation to its function in its environment.[78] To understand the microscopic origin of those macroscopic properties, "reductionism," the breaking apart of the system into its constituent parts, is a powerful tool.[79] the views are complementary and together give more understanding than either by itself.[80] This is true for any kind of system, complex adaptive or not. Looking inwards (applying reductionism) is important; but looking outwards (understanding how the system as a whole, or holistically, interacts with its environment) is also important.

Risk Assessment. The probabilistic determination of events or conditions that might prevent the maintenance or achievement of an objective. The magnitude of risk is the integral of all consequences over the probability space of their occurrence.

Robust yet fragile. "Unaffected by random component failures but vulnerable to targeted attacks on its key components"[81]

Sociotechnical systems. Systems that include both social (human) and technological elements, and that are inherently complex adaptive.

System. For our purposes, a structure organized to achieve a function.

---

[74] http://www.sandia.gov/nisac/.

[75] W. Weaver, "Science and Complexity," American Scientist **36**, 536 (1948).

[76] P. Bak, "How Nature Works: the Science of Self-Organized Criticality" Copernicus, New York (1996).

[77] D. Dennett, "Darwin's Dangerous Idea: Evolution and the Meanings of Life", Ch. 8, "Biology is Engineering" (Simon and Shuster (1995).

[78] For example: if we are using water as a heat transfer medium we are interested in its heat capacity; if we are using water as a coupling medium between rotating gears we are interested in its viscosity. The macroscopic property of the water that we care about depends on the function the water is being used for.

[79] For example, if one knows something about H and O, and how they come together to form H_2O, then one can calculate heat capacity and viscosity. And knowing this microscopic origin allows one to extrapolate to those same properties for other fluids, and to other properties that might enable the system to be used for other functions not yet thought of. Perhaps most importantly, it gives mechanistic understanding into the origin of heat capacity and viscosity, and hence for the range of environmental conditions over which the values for the properties might hold.

[80] For example, simply by "looking inwards" through reductionism one cannot know that heat capacity and viscosity were the important properties of the system to calculate for the particular function the system was being used for. One could only know that by "looking outwards" holistically at the system's function as a whole in its environment.

[81] J.C. Doyle, D.L. Alderson, L. Li, S. Low, M. Roughan, S. Shalunov, R. Tanaka, W. Willinger, "The 'robust yet fragile' nature of the Internet," Proc. National Academy of Sciences **102**, 14497-14502 (2005).

System dynamics. Developed by Jay Forrester in the 1950's at MIT as a method for enhancing learning in complex systems, to understand policy resistance and design more effective policies by considering the feedbacks and dependencies between elements of a system that influence behavior of the elements and the system as a whole.[82]

Teleonomy. The quality of apparent purposefulness and of goal-directedness of structures and functions in living organisms that derive from their evolutionary history, adaptation for reproductive success, or generally, due to the operation of a program.

TCP/IP. Transmission Control Protocol and Internet Protocol were the first two networking protocols defined. They are two of the most important and well-known protocols in an evolving suite of communications protocols. TCP/IP is actually a hierarchy of four communications (Application, Transport, Internet and Link) layers. There is a suite of protocols for each layer including those for the application layer (e.g., HTTP, SMTP, DNS, FTP) and the link layer [e.g., ARP/lnARP, Tunnels (L2TP), Media Access Control (MAC)].

Translational and bi-translational science and technology. Translational science and technology is science and technology that translates "forward" into domain-area applications or practical use. Bi-translational science and technology expands this definition to include domain-area applications that translates "backward" into developing new science or technology.[83]

Uncertainty quantification. The mathematical determination of how uncertainty in data or model structure affects the uncertainty in the results of a computational model.

Verification and validation. As applied to modeling, verification is the process of testing the model to check that it is performing as designed – both computationally and functionally; and validation is the process of evaluating the model to determine if it represents the real system sufficiently for the intended use. As with many of the other concepts presented in this paper, there is no single, accepted definition. Therefore we present a definition that is appropriate for this set of applications.

---

[82] J. Sterman, "Business Dynamics: systems thinking and modeling for a complex world," Irwin McGraw-Hill (2000); J. Forrester, Industrial Dynamics, originally published by MIT Press now available from Productivity Press, (1961).

[83] J.Y. Tsao, K.W. Boyack, M.E. Coltrin, J.G. Turnley, W.B Gauster, "Galileo's stream: a framework for understanding knowledge production," Research Policy **37**, 330-352 (2008).