Toward Agile and Resilient Large-Scale Systems:
Adaptive Robust National/International Infrastructures

S. Massoud Amin[1], and Barry M. Horowitz[2]

Keynote address at the
International Conference on Flexible Systems Management GLOGIFT-07
"Flexibility with Business Excellence in the Knowledge Economy"
November 2007, Noida, India

Abstract: How to manage or control a heterogeneous, widely dispersed, yet globally interconnected system is a serious technological problem in any case. It is even more complex and difficult to control it for optimal efficiency and maximum benefit to the ultimate consumers while still allowing all its business components to compete fairly and freely.

This paper briefly describes our on-going work in our holistic approach to analysis of the national and global infrastructure development that builds on advances in the mathematics of complexity, methods of probabilistic risk assessment, and techniques for fast computation and interactive simulation with the goal of increased agility and resilience for large-scale systems.

As an example, a model and simulation of the "Electric Enterprise" (taken in the broadest possible sense and connected to telecom, water, oil/gas and financial networks) have been developed. The model uses autonomous, adaptive agents to represent both the possible industrial components, and the corporate entities that own these components. Objectives are: 1) To develop a high-fidelity scenario-free modelling and optimization tool to use for gaining strategic insight into the operation of the deregulated power industry; 2) to show how networks of communicating and cooperating intelligent software agents can be used to adaptively manage complex distributed systems; 3) to investigate how collections of agents (agencies) can be used to buy and sell electricity and participate in the electronic marketplace; and ultimately to create self-optimizing and self-healing capabilities for the electric power grid and the interconnected critical infrastructures.

From a broader view, we have integrated these into a composite analysis technique, these advances raise an unprecedented new possibility for projecting the future implications—social, economic, environmental, human health, political, and technical—of major societal development activities and technology programs for nations individually and the world as a whole. Taken together, they promise both a real-time outlook and a future perspective on the spectrum of outcomes that might result from alternative national decision pathways. Such projection capability could reveal the development options, results, and implications for any strategy for any type of nation, whether primitive, underdeveloped, developing, or industrial. Forcing functions, critical junctures, and pinch points could be identified so that scarce development resources can be allocated to maximize benefit and minimize unintended consequences. The full realization of this next step in analysis technology will require several years of dedicated international effort, but the need is urgent and the potential payoff great. The technical—and organizational—underpinnings for such a holistic analysis approach have been demonstrated. It remains for us to build from them a global tool for a better future.

---

[1] Honeywell/H. W. Sweatt Chair in Technological Leadership, Director of Center for the Development of Technological Leadership (CDTL), Director of Graduate Studies in Management of Technology, and Professor of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55454 USA, web: umn.edu/~amin,

[2] former CEO of MITRE Corporation, Professor of Systems and Information Engineering, University of Virginia, Charlottesville, VA 22904 USA, web: sys.virginia.edu/people/bh.html

Introduction: Agile Large-Scale Systems

Around 1600, John Donne, an English poet and preacher, delivered a sermon that began: "No man is an island." Today, a less poetic, but more politically and technologically correct, version of this sentiment might be: "All human beings are interconnected through complex interactive networks and systems." For us who work with computers, the Internet and the World Wide Web come immediately to mind. Indeed, much of the world's business and industry, art and science, entertainment and crime are conducted through the Web and the Internet. But our use of these electronic information systems depends, as do our more mundane activities, on many other complex infrastructures, such as cable and wireless telecommunications, banking and finance, land, water and air transportation, gas, water and oil pipelines, and the electric power grid. All of these are, themselves, complex networks, geographically dispersed, non-linear, and interacting both among themselves and with their human owners, operators, and users.

As the complexity of these intertwined operations has increased, it has become responsible for much of "the good life" that we, at least in the more developed countries, lead today. However, with those increasing benefits come increasing risks. A common characteristic of all these networks is that local actions have the potential to create global effects by cascading throughout the network and even into other networks. These risks have not gone unnoticed. The President's Commission on Critical Infrastructure Protection issued a report in October 1997 that stressed the need for research to enhance the security of complex interactive infrastructure networks. The report cited their growing importance in many application areas, and the potentially damaging and even dangerous economic, security and health impacts of the undesirable propagation of disturbances throughout them.

There are clearly many opportunities for modeling, simulation and the use of AI in this area. Mathematical models of such systems are typically vague (or may not even exist); moreover, existing and classical methods of solution are either not available, or are not sufficiently powerful. Management of disturbances in all such networks, and prevention of undesirable cascading effects throughout and between networks, requires a basic understanding of true system dynamics, rather than mere sequences of steady-state operations. Effective, intelligent, distributed control is required that would enable parts of the networks to remain operational and even automatically re-configure in the event of local failures or even threats of failure.

Computer simulations based on the emerging science of complex adaptive systems can be used to further the understanding of the dynamics of these interactive networks. A useful approach to analyzing national and international infrastructures is to model their physical and organizational components as independent adaptive "agents" -- partly cooperating and partly competing with each other in their local operations while pursuing global goals set by a minimal supervisory function. Computer experiments would center on the investigation of agent strategies for interactions that can achieve both individual and shared goals, including competitive and cooperative strategies for system re-configuration to avoid or recover from failure. In many complex networks, the human participants themselves are both the most susceptible to failure and the most adaptable in the management of recovery. Modeling these networks, especially their economic and financial aspects, will require modeling the bounded rationality of actual human thinking, unlike that of a hypothetical "expert" human as in most applications of artificial intelligence.

What is Agility? Agility is relative; a more agile system is one that can be modified to respond to new opportunities and risks by incorporating significant new design features in a shorter period of time and in a more assured manner, than a less agile system

It is important to note that Adaptive Systems are not necessarily agile. Adaptive System is considered a system that improves its performance by dynamically adjusting to specific situations based on rules, procedures and algorithms that are built into the system. While Adaptive Systems can be designed to learn, the learning mechanisms are part of the existing system design.

Some capabilities and features of agile systems include

- Managing Time for Developing System Modifications
    - Predict future system-related opportunities and risks to provide more lead time for accomplishing design changes that are decided upon over time
        - Technical structure to reduce the required integration efforts for adding new system capabilities related to predicted opportunities and risks
        - Human organization design to more readily permit reorganizations that are related to predicted future opportunities and risks

- Managing Confidence in Rapid System Modification
    - Include fault tolerant designs that allow higher confidence in making early operational transitions of new designs
    - Provide system operators with information that increases their confidence during transition periods for new system designs

Agile Large-Scale Systems

Large mission critical systems are designed to provide assured performance, including such factors as assured capacity, reliability, response times, and failure recovery. Examples of mission critical systems where the systems engineering community has made significant contributions include our nation’s electric power system, our military’s national security systems, the international air traffic control system, and our nation’s telecommunications system.

The need for Agility stems from many systems’ needs, including:

- Reducing costs of infrastructure systems (e.g., Health Care, Power, Water, Transportation)
- Responding to business competition
- Providing responsive national security and preparedness systems
- Responding to globalization – “world systems”

Over many years, the systems engineering community has developed and refined methodologies that provide a top-down perspective to the development of such systems, including providing designers of sub-systems the design requirements that, when integrated across all relevant sub-systems, assure satisfaction of the desired integrated performance. As systems have become more and more complex, the systems engineering community has recognized that top-down system development approaches have important limitations. For example, the use of Spiral methodologies in place of Waterfall methodologies recognizes important factors in creating large systems, namely that sequential developments and operational use of parts of a large system provides an opportunity to: 1) gain value from early use, 2) learn from early use, and 3) refine overall system objectives based on actual experience with the system**.** With the continuing growth of available commercial off the shelf products, the systems engineering community has recognized that the economy of buying available components and sub-systems can warrant compromising overall desired system performance in order to avoid the development costs and lost time associated with custom development.

While the results of developing assured systems have made a significant contribution to our nation and the world, the Internet and World Wide Web provide an alternate vision for large scale systems that highlights an important shortfall in the methodologies for creating assured systems. The Internet is seen as extremely agile in responding to market forces and to opportunities afforded by new technologies. When market forces drive new technology creation and technology can be rapidly inserted into use, an extremely agile system is the end result. Of course the Internet does not provide assurances, such as response time, reliability, or security, but the market place pressures for improvement do create advances. Nonetheless, the market is not always a sufficient source for provoking assurance, as is illustrated in the situation surrounding cyber security for the Internet**,** and is also illustrated by the variation in software reliability from product vendor to product vendor, which when integrated creates the overall system's reliability. Nonetheless, the agility of the Internet, when contrasted with the very slow pace and difficulty associated with modifying assured systems is striking, and provokes the issue of what can be done to provide more agility in more assured systems.

In addition to the compelling example set by the Internet, additional factors surrounding systems engineering of assured systems need to be considered. The field of *systems biology* has been growing. Researchers in this field wish to model nature's systems so that they can better understand the opportunities for cures and for risk avoidance. Research in this field takes a bottom-up view of systems; i.e., evolution is based on survival needs of the elements (or components) of the systems, and the interdependencies derived from survival objectives of the elements create the resulting system. This is in contrast to the top-down design methodologies that systems engineers use for developing assured man-made systems**.**

As in the case of the Internet, assurances are not a direct part of modeling natural systems, but are derived from the adaptation that results from the interaction of the elements. Along with systems biology, researchers working in the area known as *complex systems* have recognized systems where there is no controlling factor and the interactions of the parts determine the performance of the resulting system, including the recognition of emergent properties; i.e., system properties that result from interactions of components as opposed to resulting from top-down design. As in the case of systems biology, adaptations by the elements of the system create the resulting system.

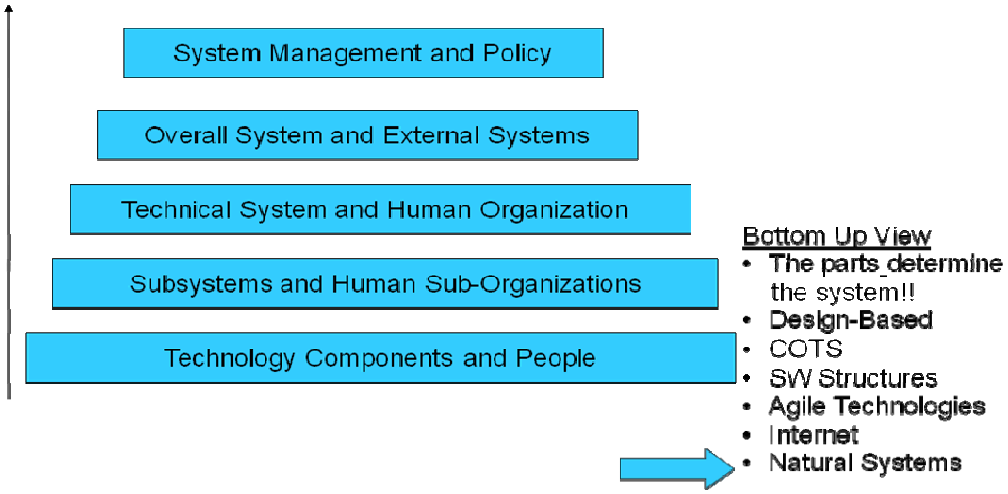

Figure 1: Multi-layered System

In particular, social scientists and economists, who have been long-time members of the man-made systems engineering community, have recognized this bottom-up view as an important aspect of recognizing human organizational and economic issues in man-made systems, and are creating various tools for forecasting the potential system level results of adaptation by the elements of the system. As indicated above we have used agent-based simulation models to forecast the adaptations of people and organizations to disruptive events.

One can expect that these research communities will continue to develop improved analysis capabilities and will carry out more and more substantial validations regarding the outcomes from their analyses. With regard to assured systems, this growing area of research suggests that some assurances, especially those tied to human and economic factors are not all that assured, and that bottom-up analysis can potentially provide an approach for helping to achieve better assurances.

In addition to the systems biology and complex system communities of research, the US Department of Defense has also recognized the relationships of adaptation as important to managing the evolution of their systems. Accordingly, the DOD has developed the concept of System of Systems to deal with emergent interdependencies and potential new supportive relationships across systems, including tri-service systems, allies’ systems, homeland defense systems, and any other pertinent systems that can improve national security.

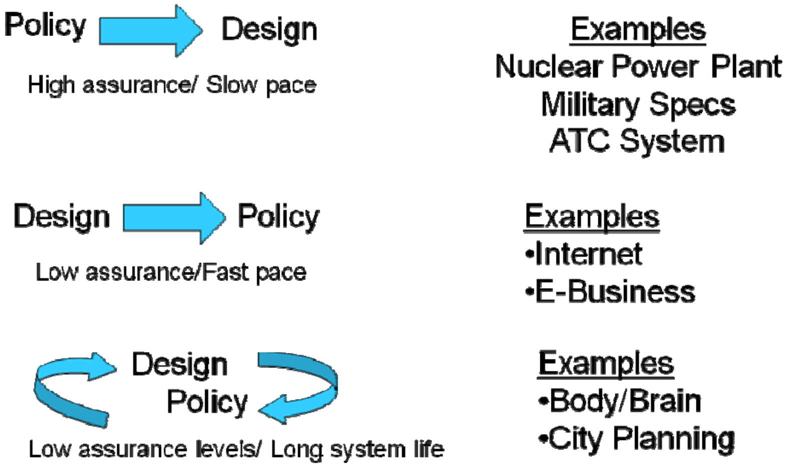

Figure 2: Policy Leads or Follows Design?

An Example: The Electricity Enterprise, Today and Tomorrow

Energy, telecommunications, transportation, and financial infrastructures are becoming increasingly interconnected, thus, posing new challenges for their secure, reliable and efficient operation. All of these infrastructures are, themselves, complex networks, geographically dispersed, non-linear, and interacting both among themselves and with their human owners, operators, and users. No single entity has complete control of these multi-scale, distributed, highly interactive networks, nor does any such entity have the ability to evaluate, monitor, and manage them in real time. In fact, the conventional mathematical methodologies that underpin today's modeling, simulation, and control paradigms are unable to handle the complexity and interconnectedness of these critical infrastructures.

Virtually every crucial economic and social function depends on the secure, reliable operation of energy, telecommunications, transportation, financial, and other infrastructures. Indeed, they have provided much of the good life that the more developed countries enjoy. However, with increased benefit has come increased risk. As these infrastructures have grown more complex to handle a variety of demands, they have become more interdependent. The Internet, computer networks, and our digital economy have increased the demand for reliable and disturbance-free electricity; banking and finance depends on the robustness of electric power, cable, and wireless telecommunications. Transportation systems, including military and commercial aircraft and land and sea vessels, depend on communication and energy networks. Links between the power grid and telecommunications and between electrical power and oil, water, and gas pipelines continue to be a lynchpin of energy supply networks. This strong interdependence means that an action in one part of one infrastructure network can rapidly create global effects by cascading throughout the same network and even into other networks.

A growing portion of the world’s business and industry, art and science, entertainment and even crime are conducted through the World Wide Web and the Internet. But the use of these electronic information systems depends, as do the more mundane activities of daily life, on many other complex infrastructures, such as cable and wireless telecommunications, banking and finance, land, water and air transportation, gas, water and oil pipelines, and the electric power grid. All of these are, themselves, complex networks, geographically dispersed, non-linear, and interacting both among themselves and with their human owners, operators, and users. Energy, telecommunications, transportation, and financial infrastructures are becoming increasingly interconnected, thus, posing new challenges for their secure and reliable operation

The North American power network may realistically be considered to be the largest machine in the world since its transmission lines connect all the electric generation and distribution on the continent. Through this network, every user, producer, distributor and broker of electricity buys and sells, competes and cooperates in an "Electric Enterprise." Every industry, every business, every store and every home is a participant, active or passive, in this continent-scale conglomerate. Over the next few years, the Electric Enterprise will undergo dramatic transformation as its key participants -- the traditional electric utilities -- respond to deregulation, competition, tightening environmental/land-use restrictions, and other global trends.

While other, more populous, countries, such as China and India, have greater potential markets, the United States is presently the largest national market for electric power. Its electric utilities have been mostly privately owned, vertically integrated and locally regulated. National regulations in areas of safety, pollution and network reliability also constrain their operations to a degree, but local regulatory bodies, mostly at the State level, have set their prices and their return on investment, and have controlled their investment decisions while protecting them from outside competition. That situation is now rapidly changing. State regulators are moving toward permitting and encouraging a competitive market in electric power.

Our research team has modeled and simulated the "Electric Enterprise" (taken in the broadest possible sense) and infrastructures connected to it, including oil/gas, water supply, telecommunications, and financial systems. The model uses autonomous, adaptive agents to represent both the possible industrial components, and the corporate entities that own these components and are now engaged in free competition. The goal in building this tool is to help these corporations evolve new business strategies for internal reorganization, external partnerships and market penetration.

Development of this tool takes advantage of recent research in Complex Adaptive Systems (CAS) which has begun to produce an understanding of complexity in natural systems as a phenomenon that emerges from the interaction of multiple, simple, but adaptive, components. Agents are no strangers to the electronic marketplace, Internet versions of this software are commonly known as "softbots" or just "bots". Most common applications have involved accessing Website contents or search engines. In contrast to earlier software, these goal-seeking agents have been semi-autonomous in achieving their objectives.

From a computer programming point-of-view, agent-based modeling and simulation is a natural extension of the prevailing object-oriented paradigm. Agents are simply *active objects* that have been defined to simulate parts of the model. Discrete event simulations with multiple quasi-autonomous agents (usually called actors or demons) have been used for at least twenty-five years to assist human decision-making in areas such as batch manufacturing, transportation, and logistics. The revolutionary new idea that comes from the computer experiments of CAS is to let the agents evolve, with each one changing in a way that adapts to its environment while that environment is modified by external forces and by the evolutionary changes in the other agents.

The agent community is allowed to evolve by causing innovative changes in the parameters of individual agents to be generated randomly and/or systematically. These parameter changes, in turn, produce changes in the agents' actions and decisions, so that the agents "tinker" with the rules and the structure of the system. Agents subjected to increased stress (resource shortages, environmental pressures, and financial losses) increase their level of tinkering until some develop strategies that relieve that stress. Some individual agents succeed (grow, reproduce, increase their profits) while others fail (shrink, die, are replaced, bought out).
Business enterprises, financial markets and the economy itself can all be viewed as complex adaptive systems and they give rise to practical problems that are often mathematically intractable. The methods developed to study CAS, as well as the insights derived from these studies, have been applied to all these areas with some success in the last decade. Practical market applications of more advanced agents represent buyers and sellers and carry out negotiations on their behalf. Agents are also used to represent stakeholders as they attempt to secure goods and services in an auction setting. Typically, the stakeholder is individual user bidding for a good. However, auctioning may not be just for individuals.

During 1998-2003 when one of the authors (Massoud Amin) was with the Electric Power Research Institute (EPRI) in Palo Alto, he funded research into agent-based auctioning as a way to address the fierce competition for resources. Electric power marketers have emerged, and wholesale electric customers are learning to shop around for the best electric suppliers. This has peaked interest in bargaining agents that trade on behalf of various stakeholders. Like agents that represent individual human users, the bargaining agents decide how much to buy, who to buy it from, how much to pay, and how they will manage the exchange of goods and money. In a power market, however, there is also concern that the entire market not be harmed by the sale. Thus, looking at how agents complete their transactions and learn from them, provide insight into the dynamics of a complex supply and demand system.

Simulations of multiple, autonomous, intelligent agents, competing and cooperating in the context of the whole system's environment have had considerable success in providing better understanding of phenomena in biology and ecology, and, more recently, in financial markets. A CAS model is particularly appropriate for any industry made up of many, geographically dispersed components that can exhibit rapid global change as a result of local actions -- a characteristic of telecommunications, transportation, banking and finance as well as gas, water and oil pipelines, and, especially, the electric power grid. The first version of this tool treats several aspects of the operation of the electric power industry in a simplified manner. For instance, it uses a DC model. However, it includes base-classes for agents representing generation units, transmission system segments, loads, and corporate owners. Users may modify and interconnect these agents through a graphical interface. Simple adaptation strategies for the agents have also been implemented. More complex ones have been designed, and implemented. Scenarios have been prepared to illustrate open access and real-time pricing.

Deregulation, Competition, Re-Regulation and New Institutions

In 1978, the United States Federal Government began the movement toward deregulation by allowing competition in several strategic sectors of the economy, starting with the airlines and followed by railroads, trucking, shipping, telecommunications, natural gas and banking. Adam Smith succinctly stated the philosophy behind this movement in 1776: “Market competition is the only form of organization, which can afford a large measure of freedom to the individual. By pursuing his own interest, he frequently promotes that of society more effectively than when he really intends to promote it.” More recently, Prof. Alfred Kahn of Cornell University, who guided the airline deregulation as the head of the Civil Aeronautics Board, expressed it in a different way: “Deregulation is an admission that no one is smart enough to create systems that can substitute for markets.”

Throughout most of the history of electric power, the institutions that furnished it have tended to be vertically integrated monopolies, each within its own geographic area. They have taken the form of government departments, quasi-government corporations or privately owned companies subjected to detailed government regulation in exchange for their monopoly status. Selling or borrowing electric power among these entities has been carried out through bilateral agreements between two utilities (most often neighbors). Such agreements have been used both for economy and for emergency backup. The gradual growth of these agreements has had the effect that larger areas made up of many independent organizations have become physically connected for their own mutual support.

In recent years, some of the local monopolies have found it beneficial to be net buyers of power from less costly producers and the latter have found this to be a profitable addition to their operations. For instance, it is typical in the western United States and Canada for surplus hydroelectric power to be transmitted south for air conditioning in the summer; while less expensive nuclear power is transmitted northward in the winter when the reservoirs are low or frozen and only nighttime heating is needed in the south. These wide area sales and the wheeling of power through non-participant transmission systems are international in extent, especially in Europe and the Americas. There is evidence of a worldwide drive to use these interconnections intentionally:

- To create competition and choice, with the hope of decreasing prices,
- To get governments out of operating, subsidizing or setting the price of electric power, and
- To create market-oriented solutions in order to deliver increases in efficiency and reductions in prices.

In order to unbundle the monopoly structure of electric power generation in the United States, Congress passed the National Energy Policy Act of 1992. National monopolies in the United Kingdom, Norway and Sweden have been de-nationalized and unbundled into separate generation, transmission and distribution/delivery companies. In most approaches to deregulation, transmission is kept as a centrally managed entity, but generation is broken into multiple independent power producers (IPP), and delivery is left to local option. New IPP are encouraged or, at least, permitted, as are load aggregators and electric power brokers, both of whom own no equipment, but are deal-makers who operate on commissions paid by the actual producers and users.

The concept behind this arrangement is that electricity, much like oil and natural gas, is a *commodity* that can be sold in the cash or spot market. As a commodity, it is possible to buy and sell future options and more complex derivative contracts based on electricity prices. However, it is not clear that electricity meets all the necessary criteria for commodity trading. The original assumptions of NYMEX and its traders were based on the model of natural gas, which, unlike electricity, can be stored economically. Once a unit of electricity is produced it must be consumed almost immediately; however, a true commodity can be stored for some length of time and consumed when and how desired. Electricity storage devices are capable of handling only a small percentage of an area's electricity requirements. Storage limitations and capacity constraints on inter-regional transfer prevent all available suppliers across the continent from head-to-head competition.

An alternative, and more entrepreneurial, view is that furnishing electricity is a *service* to the end user. Electric service may be segmented into more specific markets such as heating, cooling, lighting, building security, etc., or combined with other consumer services such as telephone, cable TV, Internet connections, etc. Both views may be reconcilable by separating the *product*, handled by generation and transmission companies, from the *service*, performed by distribution companies.

Modeling the Future

The real issue, not yet being faced in United States (or in many other nations that are moving toward greater competition in electric power) is whether such an open, competitive market can be can fair and profitable to all participants, while continuing to guarantee to the ultimate consumer of power, at the best possible price, secure, reliable electric service, of whatever quality that consumer requires.

Some utilities are contending that sudden deregulation is unfair and are seeking government reimbursement for "stranded assets" -- equipment that, for technical or financial reasons, cannot be made efficient enough to compete. In order to free the most profitable parts of their operations from regulation, other utilities are unbundling into separate and independent generation, transmission and delivery companies; or at least separate services, each optimizing its performance based on different criteria and all operating at arm's length from each other. Still other utilities are merging with, buying or being bought by, companies that may not have been in the electric power business at all. Combinations are taking place, or proposed, in which parts of former electric power monopolies join with companies whose chief product or service has been natural gas, telecommunications, cable television, engineering or finance.

Current approaches to predicting the new business structure of the electric power industry are all driven by assumed scenarios. One such scenario, based on the experience of other industries and other nations, expects that in five years there will be only a few dozen companies engaged in the actual generation, transmission and distribution of electricity. The generation companies will be completely deregulated, except for some environmental constraints. The distribution companies will still be regulated, along the lines of today's local telephone companies, but major industrial/commercial customers, and cooperatives of individual residential customers, will generate their own power or buy it from the lowest bidder. The transmission companies will be partly regulated in an attempt to ensure open access and non-discriminatory pricing for "wheeling" power between any generator and any user or distributor, while maintaining some level of system security despite their lack of control of either generation or load. However, this is just one hypothetical future scenario and various other scenarios are emerging.

The topology of these alternative scenarios/business structures dictate features of the future power system infrastructure which, in turn, suggest the most profitable re-arrangements of capital assets and market segments for each company. Hence, the predictive accuracy of this "top-down" approach depends entirely on the actual occurrence of the scenario or family of scenarios postulated. As an alternative approach, EPRI is developing a model and simulation of the "Electric Enterprise" (taken in the broadest possible sense) that uses a "bottom-up" representation of the whole system without any preconceived scenarios. Its major endogenous constraints will be the laws of physics and the cost or availability of possible technological and economic solutions. Autonomous, adaptive agents represent both the possible industrial components, and the corporate entities that own these components and are now engaged in free competition with each other. Political accommodations and corporate restructuring will appear as global emergent behavior from these locally fixed agents cooperating and/or competing among themselves. As these artificial agents evolve in a series of experiments, the simulation should expose various possible configurations that the market and the industry could take, subject to different degrees and kinds of cooperation, competition and regulation. Possible results will be the development of conditions for equilibria, strategies or regulations that destabilize the market, mutually beneficial strategies, the implications of differential information, and the conditions under which chaotic behavior might develop. This view, of course, has considerable similarity to the mathematical theory of games of strategy, but, unlike the generalized games solved by von Neumann or Nash, these are repeated games with non-zero sum payoffs. Information theoretic considerations are pertinent and these may, in turn, be represented by entropy in the state or phase space in which the system operates.

The primary goal in building this tool is to help individual companies evolve new business strategies for internal reorganization, examine the potential of entering into new partnerships or attempting to exploit new market segments. Computer experiments with this model can also provide insight into the evolution of the entire electric power industry. Within this "scenario-free" testbed, all the global behaviors that are possible in the system can emerge from local agents cooperating and/or competing among themselves in response to "what if" studies and computer experiments hypothesizing various forms of exogenous constraints. In addition, the model will serve as a practical way to estimate the benefits of implementing any proposed new technology or making hypothetical changes to existing equipment and operating practices.

Our research team is using CAS work to develop modelling, simulation, and analysis tools that may eventually make the power grid agile, resilient and self-healing, in that grid components could actually reconfigure to respond to material failures, threats or other destabilizers.

More pertinent to this conference are the areas of resilience and agility for large-scale systems that we shall present next.

Types of Large-scale Systems

There are two different types of large-scale systems, whether infrastructural or not: those with a central coordinating authority which directs and controls the design, development, and operation of the system and those that evolve through a loosely coupled arrangement among diverse partners.

The former are called simply "large-scale systems." These have the following characteristics:

- They require a long time to build – long enough so that internally and externally driven changes in desired system capabilities are very likely to occur, even while the system is under development;
- They are expensive – enough so that they must last a long time, during which it is very likely that significant new system capabilities will be required; and
- They have a large number of stakeholders (owners, users, operators) – so that over the lifetime of the system there will be a diverse set of desired changes for the system that are not likely to be universally agreed upon by the various stakeholders.

Large-scale systems that satisfy this definition include Air Traffic Control Systems, military command and control systems, and large companies' Enterprise Resource Planning (ERP) Systems, and regional Electric Power Systems.

In contrast to these centrally coordinated large-scale systems, we define "complex large-scale systems" as having the following characteristics:

- They are developed through a distributed process involving separate entities contributing in parallel to the implementation of a *shared system concept*;
- They involve organizations that each have their own objectives and risks and carry out their efforts accordingly;
- They involve a mixture of participants engaged in the development ranging from government organizations to private sector firms and perhaps to non-profit organizations;
- They do not have an integrated budget for the separate development efforts or a resource allocation plan – each entity operates on its own budget plan based on individual motivations; and

- They do not have an integrated schedule or integration plan for the separate development efforts – there are no hard start or finish dates and integration is accomplished through agreed upon cooperative experiments or in-field operational trials.

Examples of such systems would include an international environmental monitoring system, a national health care system, an intelligent highway system, an integrated intelligence system for homeland security, and the Internet.

When contrasting a *complex* large-scale system with a large-scale system, it is apparent that the uncertainties include not only those related to future opportunities and risks, but also those related to the distributed nature of its development. While the overall risks are greater, the risks of the complex large-scale system are divided among the participating organizations as opposed to being focused on the smaller number of organizations involved in large-scale system development. Also, in a complex large-scale system, new functions and capabilities can emerge in a manner that depends on the integrated contributions of individual developing organizations, possibly resulting in outcomes that were never anticipated (referred to as "emergent" properties).

Agile technologies and architectures

The ability of large-scale and complex large-scale systems designs to change in response to opportunities and risks presupposes that their human, organizational, and technical components can be changed in a timely and cost-effective manner. Current systems methodologies are not purposefully designed with such agility in mind. This thrust will focus on the conceptualization and testing of methods and processes across different domains that enable and evaluate the incorporation of agility features into systems designs. Examples of research efforts in this thrust include:

- The characterization of how Service Oriented Architecture (SOA) for software systems may be effectively deployed in large-scale systems;
- The use of standardized supply-chain management (SCM) systems to enable the rapid reconfiguration of partner firms in response to changing market conditions;
- The incorporation of wireless communication technologies into enterprise systems to foster more nimble and agile mobile workforce structures; and
- The potential for innovative organization structures – "virtual organizations" – to effectively respond to changing environmental conditions without loss of communication, coordination, or control.

Critical to the incorporation of agility features into systems is the resulting reduction in lead time from the recognition of an opportunity or risk to the system's timely response, and the cost of providing the agility features in cases where it turns out not to be needed. Trade-offs must be understood. If the lead time from measurement to system response is long, then the uncertainty of measurements will likely be large and if the cost of adding new features is large, it will discourage their incorporation into system design. Alternatively, adding agility features to a system can speed up response time, reduce the needed lead time for response, and reduce uncertainties in predictions by reducing the needed prediction window, so that change actions are more likely to be taken.

Recognizing this relationship is extremely important. It helps to explain why systems that cannot be readily changed discourage their stakeholders from exploring new opportunities and risks related to the system and can lead to very undesirable spirals that result in major lost opportunities and inability to respond to major risks that actually occur.

Agile technologies include:

- Forecasting tools for predicting the need for, and value of new capabilities
- System architectures for inserting new capabilities
- System reconfiguration tools for inserting new design features into operation
- System self-test tools for evaluating implementation of design changes
- Fault tolerance for transitioning new designs into operation

Agility designs can be developed to help reduce the pressure on forecasting:

- When many realistic opportunities and risks are possible over the same time horizon, and
- When an agile system solution (technological and organizational) can be developed to help speed up reconfiguration for all or many of the possibilities, then
- Single events no longer dominate the likelihoods for success

Agile organizational/global structures

Large-scale systems have both technical and organizational dimensions. While it is obvious that technological changes in the form of opportunities and risks will inevitably occur during and after systems design, so to will those system design elements that relate to the social elements of systems. Research efforts are needed to develop an understanding of the socio-technical aspects of large-scale systems, addressing how the organizing, staffing, and controlling elements of large-scale systems affect performance, and how they may lead to innovative structures to make systems more agile in the face of changing environmental conditions.

Among the representative areas for research investigations are:

- Development of innovative models for individual, group, and organizational behavior;
- Methods for characterizing and measuring environmental uncertainty in organizations;
- Investigating innovative structures to foster cooperative work among groups and organizations;
- Development of the easy-to-use, cost-effective methods for designing and building computational models of organizations, teams, and social systems;
- Exploration of game theoretic models to better understand cooperative and competitive behaviors in small groups, organizations, industries, and global markets;

- Investigation of how local, national, and international policies, laws, and value systems affect the design and development of large-scale infrastructure systems;
- Application of economic models, such as the Leontieff Input-Output Model, to capture and understand the effects of changes in large-scale systems as they relate to regional, national, and global competitiveness;
- Development of simulation models, in particular agent-based simulations, to further comprehend the manifold implications for designing and managing large-scale systems; and
- Field research leading to case study descriptions and analyses of large-scale and complex large-scale systems in practice.

Agile human organizations possess the following characteristics:

- Organization designed for flexible reorganization
- A culture predisposed to exploring need for important changes
- Accepted pace for making changes
- Understanding the relationships between overall system objectives and human organization to achieve desired results

Three-Response Strategies to harvest New Opportunities, Technologies and Risks:

- Responsive Top Down Design with Assurances/ Slower Response – "Assured Response"
- Bottom Up Design Changes with Time to Market Emphasis and Less Assurance – "Agility without Assurances"
- "Agility with Assurances":
  - Buying time to respond to opportunities and risks buys extra accrual of value and time for providing assurance
  - Buying time to approach "just in time" seizing of opportunities and avoidance of risks entails:
    - Early forecasting of opportunities and risks (analytical)
    - Developing and evaluating agile technology solutions (design, experimentation and simulation) across many opportunities and risks
- Buying time to approach "just in time" seizing of opportunities and avoidance of risks entails:
  - Developing and evaluating agile organizational solutions (design, experimentation and simulation) across many opportunities and risks Developing economic models for agility (costs and benefits of agile systems)
  - Developing system models and metrics for agile systems (compare alternative designs)

Figure 3 shows the overall top-down and bottom-up roles and responsibilities in the agile organizational/global structure:

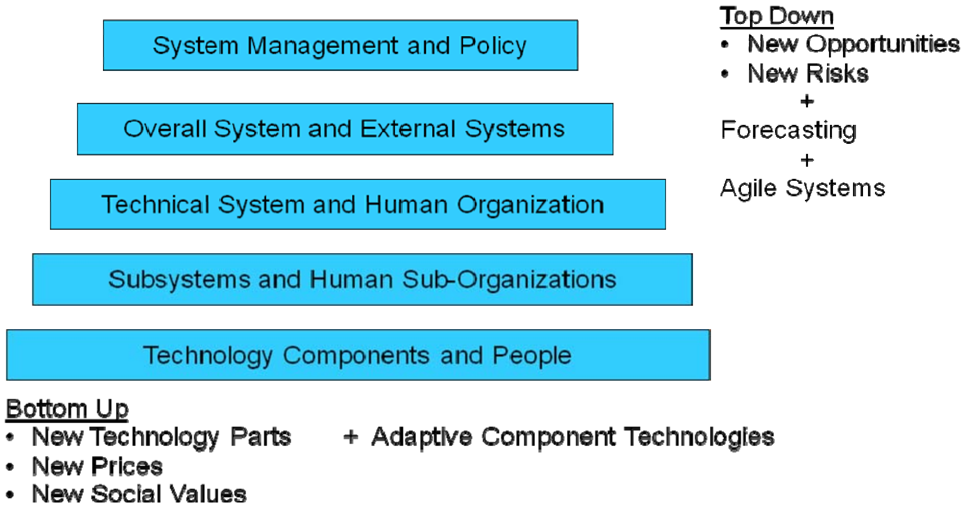

Figure 3: Agility and Multiple Layers of a System

The overall systems approach to developing agile system is depicted in Figure 4:

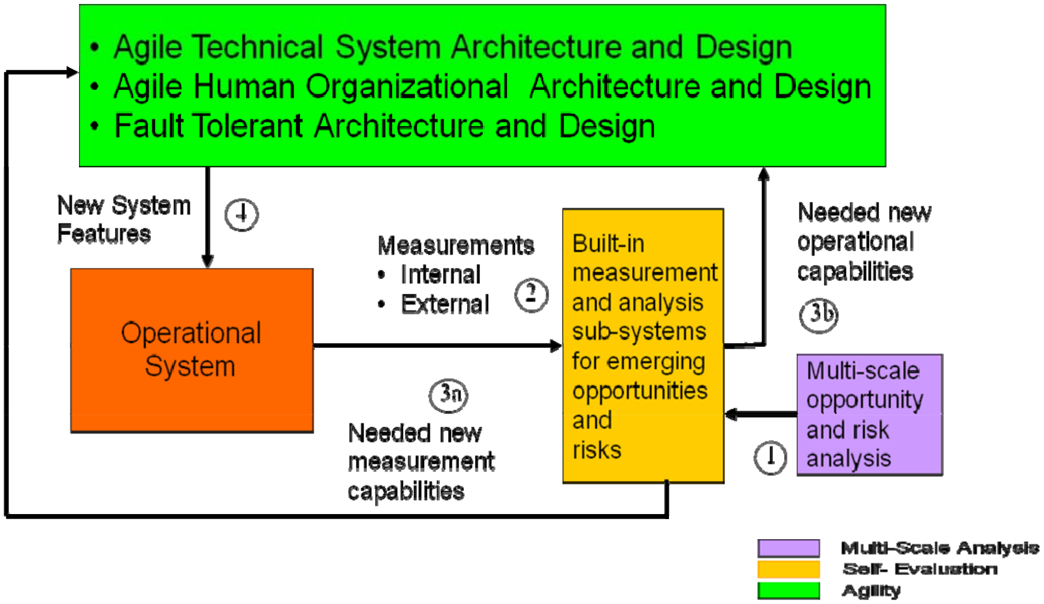

Figure 4: Design Framework for Agile Large-Scale Systems [Horowitz, 2007]

Conclusions and next steps:

There is reasonable concern that national and international, energy and information infrastructures have reached a level of complexity and interconnection which makes them particularly vulnerable to cascading outages, initiated by material failure, natural calamities, intentional attack, or human error. The electric power grid's emerging issues include creating distributed management through using distributed intelligence and sensing; use of active-control high-voltage devices; developing new business strategies for a deregulated energy market; and ensuring system stability, reliability, robustness, and efficiency in a competitive marketplace.

Managers of technically complex systems like the national infrastructure networks seldom have a detailed knowledge of underlying technologies. It is arguable as to whether or not they need such knowledge. But they do need to understand how to manage well the people who do. Only the managers can provide the internal communication network (or at least support its evolutionary development) that will give the technical people the maximum benefit of each other's information and knowledge.

Organizations, in their most abstract form, are also networks. Some of the general mathematical and computational tools for analyzing networks have found recent application in the kinds of social and communication networks that impact human performance. The well known "small world" effect has been explained as phenomena that arise only in networks that are partly ordered and partly random, and quantitative bounds have been established that characterize the degree of "information contagion" in these networks as it affects group or team decision-making.

In any particular organization, considerable benefit may be derived from knowledge developed in one part of the organization if only it can be transferred effectively to help another part. Overcoming communication barriers within organizations can be very difficult. As an example, Sencorp uses a fractal model of information, knowledge and decision-making within its large conglomerate structure to overcome these barriers.

When all else fails and there has been a general collapse of a major part of the infrastructure, emergency response may be called for. Failure of the infrastructure network may only be concomitant to a natural disaster, like an earthquake or flood. But it is important that emergency response actions not further degrade the telecommunications, electric power, gas and water pipelines, etc., all of which can be of great assistance in mitigating the effects of the disaster. The usefulness of computer simulation has been demonstrated for both the analysis of emergency situations and for training people in their management. These simulations provide realistic representations in real-time of the damage caused by various natural and human-caused disasters. They can also provide statistical estimates of the results of human attempts to control this damage, for instance, the effects of a selected method for fighting a fire. But often, the greatest source of uncertainty in the management of emergencies is the behavior of the people directly affected by it. The training of emergency response teams emphasizes systematic, dependable performance even, if necessary, at the expense of rapidity and flexibility. Such dependability can come only after many repetitions of the training, to the point where details of behavior are almost automatic. This level of dependability is essential so that the plan for emergency management can safely assume that the team will perform as required. However, the behavior of the general population directly affected by the emergency is much less dependable. Emergency managers have little basis for predicting it, and only limited ability to control it.

Past experience with other emergencies of a similar kind, sometimes even in the same geographic area, can provide rough guidelines. A time history including sequences of equipment failure and patterns of human response can now be assembled from the modern instrumentation being installed in most infrastructure networks as well as from direct reports by volunteers calling in on cellular phones. Using these collected space-time histories to forecast the behavior of the affected individuals may be possible using hierarchical clustering.

Any complex dynamic infrastructure network typically has many layers, decision-making units and is vulnerable to various types of disturbances. Effective, intelligent, distributed control is required that would enable parts of the networks to remain operational and even automatically re-configure in the event of local failures or threats of failure. In any situation subject to rapid changes, *completely centralized control requires multiple, high-data-rate, two-way, communication links, a powerful central computing facility, and an elaborate operations control center. But all of these are liable to disruption at the very time when they are most needed,* i.e., when the system is stressed by natural disasters, purposeful attack, or unusually high demands. Management of disturbances in all such networks, and prevention of cascading effects throughout and between networks, requires a basic understanding of the true system dynamics, as well as effective distributed control functions to enable parts of the networks to remain operational or even to automatically re-configure themselves in the event of a threat or other potentially destabilizing disturbance.

When failures occur at various locations in such a network, the whole system breaks into isolated "islands," each of which must then fend for itself. With the intelligence distributed, and the components acting as independent agents, those in each island have the ability to re-organize themselves and make efficient use of whatever local resources remain to them in ways consonant with the established global goals to minimize impact on the overall network. Local controllers will guide the isolated areas to operate independently, while preparing them to rejoin the network without creating unacceptable local conditions either during or after the transition. A network of local controllers can act as a parallel, distributed computer, communicating via microwaves, optical cables, or the power lines themselves, and intelligently limiting their messages to just that information necessary to achieve global optimization and to facilitate recovery after failure.

From a broader view, the routine of an integrative approach for holistic technological/economic/social outcomes should be expected as another important byproduct of this effort. The effect in academia would be to broaden the scope of teaching in the applied scientific and engineering specialties, as a broad mode of thought becomes commonplace and the technology faculties become involved in fundamental R&D.

Moreover, the sophistication of international academic and political participants would be enhanced. Hopefully international collaboration will become a habit. However it should be recognized that the global problems being addressed will take a long time to dissect, and partial progress should be taken as encouraging.

It is the authors' firm belief that technology development and its effective management is an important key to improving the quality of life for coming generations, a key that adds hope for the planet and for those who will inherit the future. In the coming century there will be many new choices made about technology development, and it will be important for decision makers to have a systematic, reasonably defensible, and comprehensive perception of the future social impact and feasibility of these choices in order to make wise choices and invest our wealth well. The projection tool that is under development provides an opportunity for the "have" nations to use their technology capability and resources to help the "have not" nations identify for themselves an appropriate level of development and economic activity.

While the immediate and critical goal of work has been to avoid widespread network failure, but the longer-term vision is to enable adaptive and robust infrastructure; as expressed in the July 2001 issue of the Wired magazine: *"The best minds in electricity R&D have a plan: Every node in the power network of the future will be awake, responsive, adaptive, price-smart, eco-sensitive, real-time, flexible, humming - and interconnected with everything else."*

Achieving this vision and sustaining infrastructure reliability, robustness, resilience, agility and efficiency are critical long-term issues that require strategic investments in research and development. Much research remains to be done ranging from mathematical underpinnings and foundational work (Figure 5) to applications to organizational and global agility. Given economic, societal, and quality-of-life issues and the ever-increasing interactions and interdependencies among infrastructures, this objective offers exciting scientific and technological challenges.

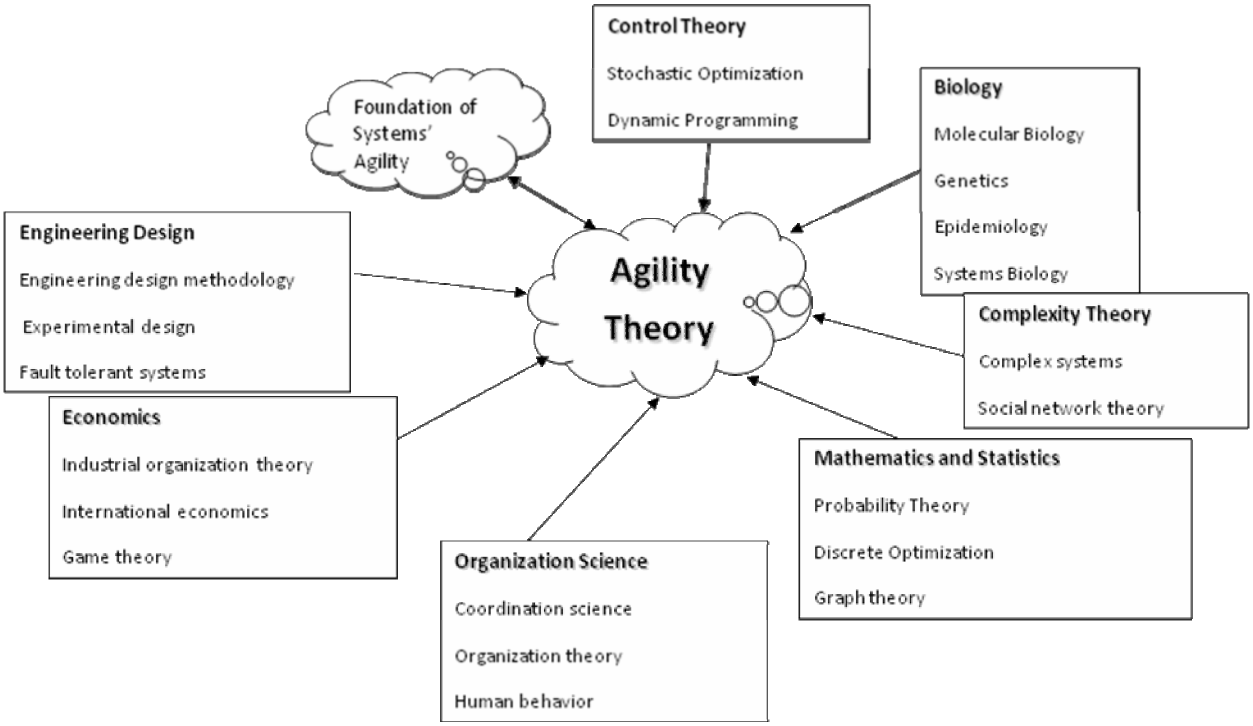

Figure 5: Overview and Taxonomy of disciplines incorporated within "Agility Theory," however foundation remain fragmented and there is a need for mathematical formalism

References

1. Amin, M., and P. Schewe, "Preventing Blackouts," Scientific American, pp. 60-67, May 2007

2. Amin, M., "New Directions in Understanding Systemic Risk", with NAS and FRBNY Committee, National Academy of Sciences and Federal Reserve Bank of NY, 77 pages, March 2007

3. Amin, M., and B.F. Wollenberg, "Toward a Smart Grid," IEEE Power and Energy Magazine, Vol.3, No 5, pp. 34-38, Sept/Oct. 2005

4. Amin, M., “Scanning the Issue,” *Special Issue of Proceedings of the IEEE*, Vol. 93, Number 5, pp. 855-860, May 2005

5. Amin, M., “Energy Infrastructure Defense Systems,” *Special Issue of Proceedings of the IEEE*, (Amin), Vol. 93, Number 5, pp. 861-875, May 2005

6. Amin, M, Special issues of IEEE Control Systems Magazine on Control of Complex Networks, Vol. 21, No. 6, December 2001 and Vol. 22, No. 1, February 2002

7. Amin, M., "Powering the 21st century: We can—and must—modernize the grid," (Amin), IEEE Power and Energy Magazine, pp. 93-96, March/April 2005

8. Amin, M., “Electricity,” in *Digital Infrastructures: Enabling Civil and Environmental Systems through Information Technology*, R. Zimmerman and T. Horan (Editors), Chapter 7, pp. 116-140, July 2004

9. Amin, M., "North American Electricity Infrastructure: System Security, Quality, Reliability, Availability, and Efficiency Challenges and their Societal Impacts," Chapter 2 in the National Science Foundation (NSF) report on "Continuing Crises in National Transmission Infrastructure: Impacts and Options for Modernization," June 2004

10. Amin, M., “Balancing Market Priorities with Security Issues: Interconnected System Operations and Control under the Restructured Electricity Enterprise,” IEEE Power and Energy Magazine, Vol.2, No 4, pp. 30-38, July/August 2004

11. Amin, M., “North American Electricity Infrastructure; Are We Ready for More Perfect Storms?” *IEEE Security and Privacy* magazine, volume 1, no. 5, pp. 19-25, Sept./Oct. 2003

12. Amin, M., “Restructuring the Electric Enterprise: Simulating the Evolution of the Electric Power Industry with Adaptive Agents,” Chapter 3, pp. 27-50, in *Electricity Pricing in Transition*, Ahmad Faruqui and Kelly Eakin (editors), Kluwer Academic Publishers, September 2002

13. Amin, M., “Security Challenges for the Electricity Infrastructure,” Special issue of the *IEEE Computer Magazine* on Security and Privacy, April 2002

14. Amin, M., and Lee, S. "Evolving Energy Enterprise: Possible road ahead and challenges for R&D," Proceedings of the International Conference on Electrical Engineering, panel on Future development of power systems in the 21st century, paper # 0324, 8pp., Xi'an, China, July 22-26, 2001

15. Amin, M., “Toward Self-Healing Energy Infrastructure Systems,” cover feature in the *IEEE Computer Applications in Power*, pp. 20-28, Vol. 14, No. 1, January 2001

16. Amin, M., and Ballard, D. “Defining New Markets for Intelligent Agents,” *IEEE IT Professional*, pp. 29-35, Vol. 2, No. 4, July/Aug. 2000

17. Amin, M., “Modeling and Control of Electric Power Systems and Markets,” *IEEE Control Systems Magazine*, pp. 20-25, Vol. 20, No. 4, Aug. 2000

18. Amin, M., “Toward Self-Healing Infrastructure Systems,” cover feature in the *IEEE Computer Magazine*, pp. 44-53, Vol. 33, No. 8, Aug. 2000

19. Amin, M., “National Infrastructures as Complex Interactive Networks,” chapter 14 in: *Automation, Control, and Complexity: An Integrated Approach*, T. Samad & J. Weyrauch (Eds.), pp. 263-286, John Wiley and Sons Ltd., NY, March 2000

20. Amin, M., and Wildberger, M, "Human Performance Issues in the Security of the National Infrastructure," 6 pp., Proceedings of the ASTC2000 (Adv. Simulation Technology Conf.), Washington, DC, April 16-20, 2000

21. Amin, M., and Holmes, R. "Adaptive Infrastructures – the Complex Interactive Networks and Systems Initiative," 9 pp., Proceedings of the ISA 2000 conference, San Antonio, TX, June5-9, 2000

22. Bak, P., 1996. *How Nature Works: The Science of Self-Organized Criticality,* Springer-Verlag, New York, NY.

23. Bar-Yam, Yaneer, "About Engineering Complex Systems: Multiscale Analysis and Evolutionary Engineering," in: Engineering Self Organising Systems: Methodologies and Applications, Brueckner, S., G. Di Marzo Serugendo, A. Karageorgos, R. Nagpal ,Eds., ESOA 2004, LNCS 3464, Springer-Verlag, 2005, pp. 16-31.

24. Boehm, Barry, "A Spiral Model of Software Development and Enhancement," Computer, May 1988, pp. 61-72.

25. Boehm, Barry and Richard Turner, "Management Challenges to Implementing Agile Processes in Traditional Development Organizations," IEEE Software, September-October 2005, pp. 30-39.

26. Bonifacio, M., R. Cuel., G. Mameli, And M. Nori, "A peer-to-peer architecture for distributed knowledge management," in: Proceedings of the 3rd International Symposium on Multi-Agent Systems, Large Complex Systems, and E-Businesses, 2002.

27. Buede, Dennis M., The Engineering Design of Systems: Models and Methods, New York: John Wiley & Sons, 1999.

28. Chadwick, D. J., V. M. Patel, and L. G. Saxton, "Communications Architecture for Early Implementation of Intelligent Vehicle Highway Systems," Transportation Research Board, 1993.

29. Cowan, G.A., D. Pines, and D. Meltzer. Eds. 1994. *Complexity: Metaphors, Models, and Reality*. Addison-Wesley: Reading, MA.

30. de Neufville, Richard, "Uncertainty Management For Engineering Systems Planning And Design," Engineering Systems Monograph, MIT Engineering Systems Division, March 29-31, 2004.

31. EPRI. 1999. "Simulator for Electric Power Industry Agents (SEPIA): Complex Adaptive Strategies", November 1999, Palo Alto, CA, Report TR-112816

32. EPRI. 1999. "Prototype Intelligent Software Agents for Trading Electricity: Competitive/ Cooperative Power Scheduling in an Electronic Marketplace", Dec. 1999, Palo Alto, CA, Report TR-113366

33. EPRI. 2000. "E-Commerce Applications and Issues for the Power Industry", May 2000, Palo Alto, CA, Report TP-114659

34. Friedman, D. 1991. "Evolutionary Games in Economics." in *Econometrica* 59,3 (May): 637-666.

35. Haimes, Y. Y., N. C. Matalas, J. H. Lambert, B. A. Jackson, and J. F. R. Fellows, "Reducing the vulnerability of water supply systems to attack," Journal of Infrastructure Systems, American Society of Civil Engineers, 4(4): 164–177, 1997.

36. Haimes, Y. Y., Risk Modeling Assessment, and Management. New York: John Wiley & Sons, 1998.

37. Harp, Steven A., Sergio Brignone, Bruce Wollenberg, and Tariq Samad, "SEPIA: A Simulator for Electric Power Industry Agents", IEEE Control Systems Magazine, pp. 53-69, Vol. 20, No. 4, August 2000

38. Holland, J. H. 1995. *Hidden Order: How Adaptation Builds Complexity*. Addison Wesley, Reading, MA.

39. Horowitz, B. M. “Self-Evaluating Agile Large-Scale Systems (SEALS),” Proceedings of the Symposium on Complex Systems Engineering, California State University, Los Angles, January 11-12 2007. Available at http://cs.calstatela.edu/wiki/images/d/d3/Horowitz.doc

40. Hughes, T. P., “Rescuing Prometheus, Four Monumental Projects That Changed the Modern World,” Pantheon Books, 1998

41. Leppingwell, John W.R., “Soviet Strategic Air Defense and the Stealth Challenge,” International Security, Fall 1989, Vol. 14, No. 2, pp. 64-100.

42. Reich, Yoram, et al., “Building Agility for Developing Agile Design Information Systems,” Research in Engineering Design, Vol. 11, No. 2, August 1999, pp. 67-83.

43. Sharma, Chetan, Wireless Internet Enterprise Application, New York: John Wiley & Sons, 2001.

44. Soros, G. 1995. *Soros on Soros: Staying Ahead of the Curve*. John Wiley & Sons, Inc., NY, NY.

45. Starr, C., and M. Amin, “Global Transition Dynamics: Unfolding the Full Social Implications of National Decision Pathways,” 11 pp., submitted to the President of the US National Academy of Engineering, September 2003

46. Sutherland, J. and W. van den Heuvel, “Enterprise application integration and complex adaptive systems,” Communications of the ACM, Vol. 45, No. 10, October 2002, pp. 59-64.

47. Wang, Tao and Richard de Neufville, “Identification of Real Options ‘in’ Projects,” Proceedings of the 16th Annual International Symposium of the International Council on Systems Engineering (INCOSE), Orlando, Florida, July 2006.

48. Zhao, T., and Tseng C. “Valuing flexibility in infrastructure expansion.” Journal of Infrastructure Systems, Vol. 9, No. 3, September 2003 pp. 89-97.