Structure and Dynamics of Multiple Agents in Homeland Security Risk Management

Gary G. Nelson
Homeland Security Institute
Arlington, VA
gary.nelson@hsi.dhs.gov

ABSTRACT

Homeland security risk-management invokes extreme problems of uncertainty and resource allocation. But there is also an evident evolutionary approach to the problem, since every complex system we see is a successful competitor for limited resources under a large spectrum of contingencies. It is vital to discern the structural and dynamical principles of such evolutionary systems, extended from the strict biological process to human organizations and artifacts. These principles can be used to correct pathologies of human risk-manageing organizations and then to exploit recent technologies of artificial-agent ecologies and open-communication systems in hybrid human/artificial risk-management systems. This paper: 1) Summarizes the structural and dynamic principles of complex adaptive systems (CAS) in risk-management terms; 2) Examines "resilience" as the descriptor of risk-managing CAS, and; 3) Applies hybrid human/artificial multi-agent systems (MAS) in a risk-managing, resilient, network.

Introduction: Evidence of Evolutionary Risk Management

The problem of homeland security (HLS) is to allocate scare resources for prevention, protection and response against a spectrum of uncertain and severe hazards. The systems of attack and security are both complex social, economic and political systems within—since natural hazards must not be forgotten—a natural environment. The long-horizon discounting and quasi-equilibrium of the market are not equal to extraordinary hazards and even public resources can support only limited preparedness. The limits to this are in our own social system. As stated by the Secretary of the Department of Homeland Security (DHS), "Our goal is to maximize our security, but not security 'at any price'" [Chertoff, 2005].

Extraordinary risk management is a hard problem. Yet everything we see is a persistent complex-adaptive system (CAS) that has evolved *by competing for resources under a large spectrum of contingencies*. The complexity of society—that distributes decisions in a multi-agent system (MAS) beyond any objective analysis or central planning—provides the "resiliency" against a spectrum of hazards that exceeds equilibrium (or statistically stationary) bounds. Put another way, there is an empirical

CAS/MAS solution to the risk management problem where there is no apparent analytical, central-planning solution based on economic metrics and probabilisties.

Experience indicates that human structural and conceptual problems in large-scale planning generally constrain the application of CAS/MAS. Transportation planning has been one of our most ambitious social endeavors. Yet we use the ongoing "transportation crisis" not as an indication of bad risk management but instead to justify "more of same" [Nelson and Allen, 2002]. The vision for the next-generation air traffic management (ATM) system [RTCA, 2002] features the use of "intelligent agents" with ubiquitous data and communications. However, decades of applying computation and communications technologies have not much changed old, manual, ATM practices. There has been over a half century of codifying systems engineering for the development of complex systems, which is a risk management activity. Yet organizations still have spectacular failures in such developments despite active work to apply the lessons of CAS to this field [Norman and Kuras, 2004, Nelson, 2006a]. CAS/MAS approaches should apply as the architecture for systems engineering and risk management in HLS [Nelson, 2006b].

Risk Management

Risk management (RM) is defined here as multi-agent contention for resources in an "environment". Each agent experiences that environment as a physical (geographical) and logical (informational) state-space trajectory through time. The environment is the set of all the paths of agents that contend purposefully plus "nature". Every agent wants to predict and affect the environment for a resource payoff. The "risk" is the predictive uncertainty about payoff. That uncertainty is relative to the information generated as the network of space-time trajectories of all agents that is the CAS/MAS.

The essential problem of RM is prediction of the state of the environment and there are two dynamical cases for this: The equilibrium, or statistically stationary case in which a predictive model can be optimized once and for all, and the non-equilibrium or statistically nonstationary case that typifies CAS. The key dynamic parameter of prediction is scale: This can be related to information samples over spatial (ensemble) and temporal (time-series) dimensions.

The scaled-prediction network model of agent paths ignores the conventional distinction between games among purposeful agents, and "games against nature". HLS intends to be "all hazards" and should ignore the unfortunate dichotomy between RM of terrorism (intentional attacks) and accidents or natural hazards. Both purposeful and natural agents are processes that generate information in a nonequilibrium environment. The environment consists of scaled collectives of agents, both human and natural with corresponding scales of predictability and controllability. RM requires as much useful predictive information as possible, leading to principles about information sharing and collaboration, even with agents representing nature *in* the social environment (e.g., weather forecasters).

Summary of CAS Principles

Definition: CAS are systems of multiple agents that are organized as scale hierarchies to span the nonstationarity of information-generation by agents pursuing payoffs (contending over resources) in the physical and logical space-time of the environment.

Structure: The Scale Hierarchy

There are two branches to hierarchy theory: Specification hierarchies (categorical set containments) and scale hierarchy (dynamical process containments). Scale hierarchy theory has been developed by Salthe [1985, 1993, 2001]. The present author, through dialog with Salthe, has shown the application of the model in air traffic management (ATM) [Nelson, 1992, 2003] and the relation to the kinds of stochastic processes and prediction in such systems and particularly weather [Nelson, 1990]. The structure has been applied as a principle in the architecture of HLS [Nelson, 2005]. The RM scale-hierarchy model consists of portfolios as collectives of peer agents for RM at any one scale. These agents *represent* (argue on behalf of for budget resources) potential payoff paths under *various* environments. Diversity of the agents (hedging) is essential to the RM portfolio. The scale hierarchy means that a portfolio can commit different *aggregations* of resources but leaving a detailed degree of freedom to lower scales.

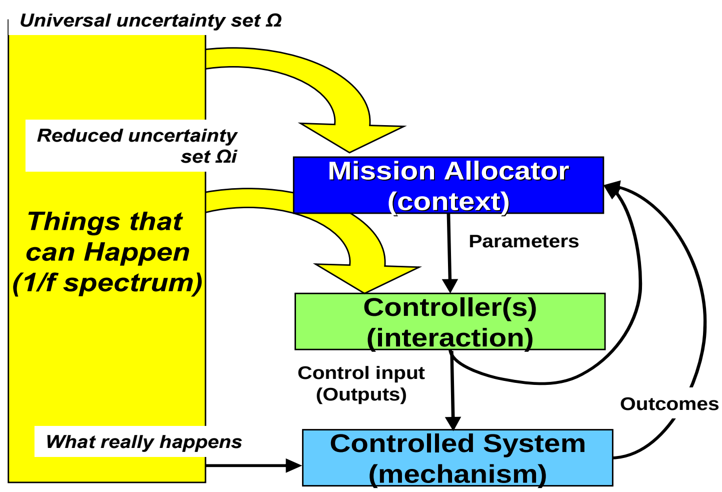

Figure 1: The Scale Hierarchy with Learning and its Uncertainty Environment

Figure 1 shows the scale hierarchy in the form of self-organizing control [Mesarovic, 1964, Lendaris, 1964] with learning in a stochastic environment. The learning scheme generally follows reinforcement learning (RL) but also represents the Government Performance Results Act (GPRA) that is one of the budgeting and acquisition reference models of the Office of Management and Budget (OMB). The absolute number of scales must span the uncertainty (information generation) of the environment. The table below gives three interpretations of this triadic structure

Table 1: Interpretations of the Scale Hierarchy Triad

Triadic Scale	Adaptive Control Model	Learning/GPRA Model
Top (context, goal)	Allocates parameters that define the problem to be solved	Compares outcomes with goals; adjusts penalty/reward to agents in the focal network.
Middle (focal, interacting)	Decides control inputs to a controller (what state to achieve as problem solution)	Strives to maximize a local utility based on contention with peers; produces a decision as output (information).
Bottom (mechanism, instantiation of environment as multiple local states)	Converts the real and local environment into an outcome (new local-environment state)	Represents the environmental actuation (outcome) of the decision.

Scale Hierarchy Dynamics

A general characterization of the uncertainty environment applies to CAS/MAS. Figure 1 on the left shows a spectrum of "things that can happen" in any space-time sample of the environment. Figure 2 formalizes this characterization as a *power spectrum* and a scaled behavior of its statistics.

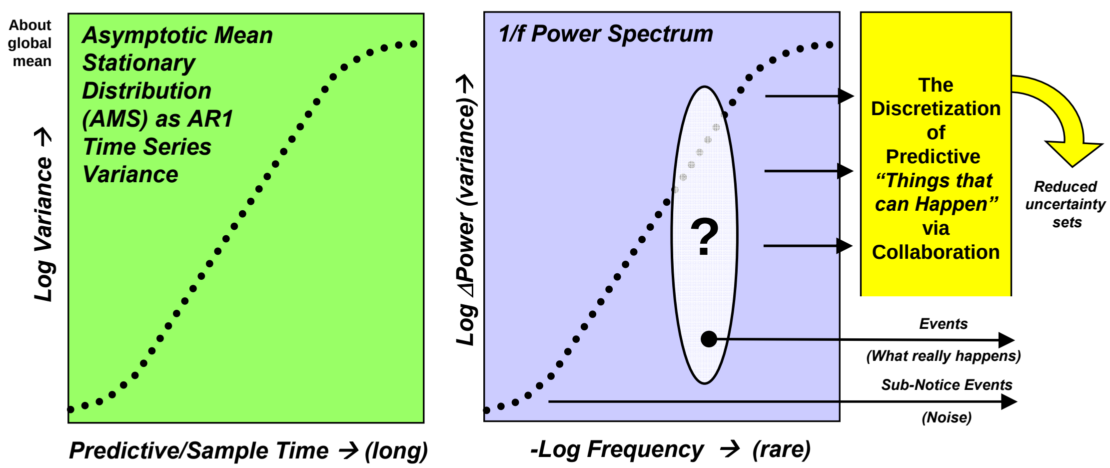

Figure 2: AMS Processes and the 1/f spectrum

The 1/f power spectrum is known for natural disasters and accidents, is the basis for risk management tools such as CORA®, and has been identified in intentional terror events [Clauset and Young, 2005]. Strictly, the power spectrum truncates at its extremes so the process has finite power over ensembles or time series. Statistics of any of these samples will be nonstationary as described by the asymptotic mean stationary (AMS) process [Gardner, 1990] at the left of figure 2. The AMS process says that you can decide locally but not over an indefinite scale unless you are willing to accept exponentially increasing error with scale. The scale hierarchy of portfolios makes the tradeoff between a local/reliable prediction of payoff (efficiency), and a grosser but more robust prediction (resiliency). There are appropriate portfolios of decisions for each scale that encompass increasing aggregations of social organization and resources.

The Mesarovic scheme shows the "universal uncertainty set" Ω that is decomposed and allocated to lower scales. In the HLS case, the national scale of governance (policy) has indeed allocated a set of "national planning scenarios" of hazards that stipulate more detailed (administrative) RM planning [HSC, 2005]. Allocation of exponentially divided uncertainty to portfolios (clusters of agents) at a lower scale reflects the fundamental principle of partial-decomposability of Simon and Ando [1961].

Multi-Agent Systems for Risk Management

RM requires the application of innovation within an open architecture. The hybrid CAS/MAS supposes a progression, as shown in the three views of figure 3:

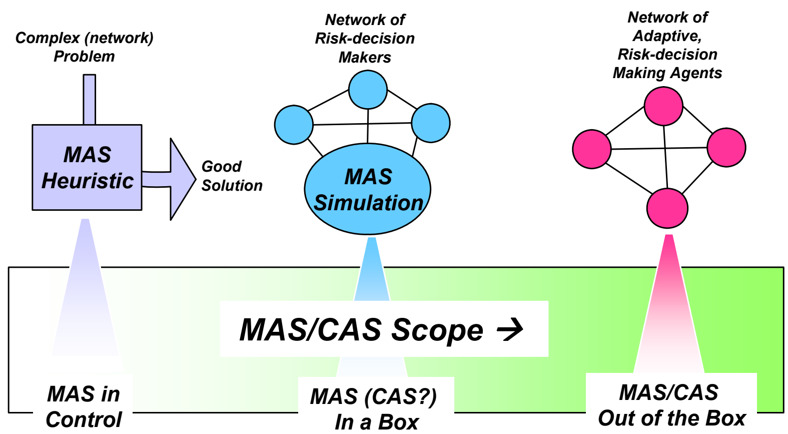

Figure 3: A Typology of MAS/CAS

On the left is an optimizing (or heuristic) algorithm as MAS-in-a-box within a fixed, human-imposed, problem context. In the middle the MAS interacts with a network of human agents in resource contention. This simply makes explicit the interaction scale that alters the context for the MAS-in-a-box. On the right is a contentious network of agents without distinction between human and artificial. The MAS of multiple human and artificial agents *is* a CAS.

The Possibility of Resilience

Resilience is a desirable attribute for a HLS RM system, and some often-cited definitions of resilience are:

1. Amount of change a system can undergo and still control its function and structure. The ability to self-organize, to learn. [Holling, 1973]
2. Expected number of network node pairs that remain connected. [Colbourn, 1987]
3. Largest number of nodes failed (with probability p) such that network remains connected with probability 1-p. [Najjar and Gaudiot, 1990]
4. k-edge/node failure resilient: If failures ≤ k, each subnetwork is self-sufficient. [Rosenkrantz et. al. 2004]

Figure 4 shows an architecture of layers involved in HLS RM.

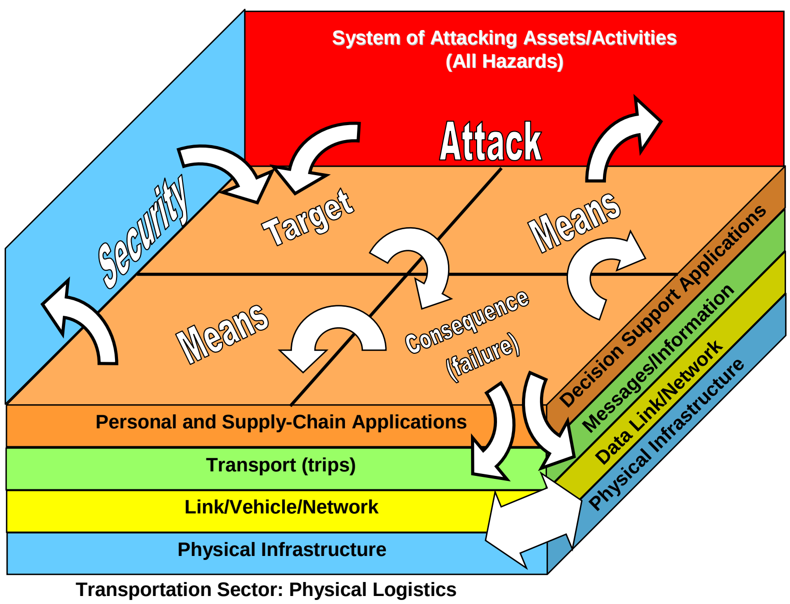

Figure 4: An Information/Physical-Interactive Architecture of Security

Resilience must apply to all the infrastructures that support the logical RM network. These infrastructures are networks of physical transportation or informational communications. All CAS are dissipative so a physical energy and material flux is assumed, and may be placed within "transportation". The Hollings definition is most appropriate to a CAS, but the old metrics of reliability, maintainability and availability (RMA) still apply to define the complex RM problem in another way:

- Reliability: Which of many, many assets under the control of many jurisdictions do we want to make more reliable under a set of hazards of indefinite potency and ill-defined frequency of occurrence?

- Maintainability: What assets do we establish (and themselves maintain) so that failures of the various assets are reduced to an "acceptable" duration? This also raises questions of whether there are lesser levels of functionality to be achieved in various times after failure.
- Availability: What are the network paths and types of functionality (services) that are substitutable, or able to be foregone, by who and for how long, in order to achieve an acceptable social outcome?

Making the Most of Agents in and for Networks

The definition of RM as a network problem about the resiliency of the problem-solving (and agent-sustaining) network is consistent with the evolution of a hybrid CAS/MAS. The real system in figure 4 answers the complex RMA questions by resilient scaled portfolios of agents prospectively finding paths through the contingent environment This suggests developing RM as a congested network problem. The most general version of the capacitated-resource problem is the El Farol Bar problem [Arthur, 1994 and Challet et. al., 2004].

In the case of flow-network assignment problems, Wardrop's Principles [1952] long ago defined the divergence between a user-equilibrium assignment (where user =agent) and a system-optimal assignment. This divergence applies to the "wicked" problem [cf. Rittel and Webber, 1973] of social resource allocation in RM that obviates implicit conditions for any centrally-planned "system optimum":

- The goal cannot be reduced to money-metrics with collective acceptablity.
- There is no pareto-optimal solution, but rather uneven distributional effects.
- The uncertainties (network paths to failure, modes of attack, costs of security) are great and lack any credible aggregation into statistics.
- We do not even know where a "more satisficing" solution is in the vast collective state space.

The scale hierarchy is motivated where local agent-path satisficing diverges from evident collective goals. The so-called Braess' paradox [1969, Tumer and Wolpert, 2000] illustrates such divergence remedied through *removal* of some link resource. A form of congestion leading to divergence in security is the Interdependent Security (IDS) problem [Kunreuther and Heal, 2003]. The reality of the scale-hierarchy of social governance, and the CAS/MAS principles reconcile such divergence with the vertical-scale feedbacks implemented as RL in Collective Intelligence (COIN) [Wolpert et. al, 1999 etc.]. COIN is fertile for extension to the CAS model at the right of figure 3.

Conclusion

The hybrid CAS/MAS can be applied to a properly stated homeland security risk-management (HLS RM) problem, and to any "wicked" collective RM problem. The network-assignment of agent paths must move off some quasi-equilibria to others along with evident shifts in the “environment” regarding terror, internal social problems and the natural environment. This is a collective contention/collaboration problem with urgency in the face of exponential human population growth and physical-resource consumption that will enter the far reaches of the contingency spectrum. Security is just one set of dimensions of the payoff and we should not overly “stovepipe” the problem.

A model has been provided for hybrid CAS/MAS to facilitate the solution of portfolio “path finding” problems by a scale hierarchy by of fast/dense artificial agents that multiply human physical capabilities. But our human organizations often diverge from the CAS/MAS model of RM and resilience and this is the barrier to its hybrid extension: We use people as well as automata as “MAS in a box” in the well-known “mushroom” syndrome. Whether for humans or automata, we partition into stovepipes and then forget about the essential interoperability of peers. The context scale turns this falure of coordination into a central-palnning problem that “puts Humpty Dumpty together again” by substituting analysis for evolution. Continued citation of the information sharing and collaboration problem is symptomatic of this. Worse, higher scales tend to defer problems to lower scales without properly solving *their* allocation problem in the first place. Let us not put the technical opportunities handed to us as new wine in old jugs of stovepiped, poorly allocated, manual functions. If we can see how to organize humans properly for the RM problem, then the hybrid will evolve.

References

Braess D., 1969, Über ein Paradoxon aus der Verkehrsplanung. Unternehmensforschung 12, 258 – 268

Arthur, W. Brian, 1994 (updated 2006). Inductive Reasoning and Bounded Rationality (The El Farol Problem). Available at http://www.santafe.edu/arthur/Papers/El_Farol.html

Challet, Damien, Matteo Marsili, and Yi-Cheng Zhang, 2004. Minority Games: Interacting agents in financial markets, Oxford University Press.

Chertoff, Michael, July 13, 2005, remarks as prepared at http://www.dhs.gov/dhspublic/display?content=4597

Clauset and Young, 2005. Scale Invariance in Global Terrorism, at http://www.arxiv.org/PS_cache/physics/pdf/0502/0502014.pdf

Colbourn, C.J., 1987, Network Resilience, SIAM J. Alg.

Gardner, William A., 1990. Introduction to Random Processes: With Applications Signals and Systems. Second Edition, McGraw Hill.

Holling, C.S., 1973, Resilience and stability of ecological systems, Annual rev. ecological systems.

HSC: Homeland Security Council 2005. National Planning Scenarios, Version 20.2, April 2005. (For Official Use Only). An older public version may be seen at http://www.globalsecurity.org/security/library/report/2004/hsc-planning-scenarios-jul04_intro.htm

Kunreuther, Howard and Geoffrey Heal, 2003, Interdependent Security, J. of Risk and Uncertainty, 26(2-3), pp. 231-249.

Lendaris, George G. 1964. On the Definition of Self-Organizing Systems, Proc. Of the IEEE, March 1964.

Mankins, John C. 1995. TECHNOLOGY READINESS LEVELS, A White Paper, April 6, 1995, Advanced Concepts Office, Office of Space Access and Technology, NASA. At http://www.hq.nasa.gov/office/codeq/trl/trl.pdf

Mesarovic, M.D., 1964. Self-Organizing Control Systems, pp. 265-269, IEEE Trans. On Applications and Industry, 83.

Najjar, W. and J-L Gaudiot, 1990, Network Resilience, IEEE Trans. Computing.

Nelson, Gary G., 1990. Hierarchies, Noise Like the Wind and Networks, pp. 564-570, Vol. II, Proceedings of the 34th Annual Meeting, International Society for the Systems Sciences, Portland, OR. July 8-13, 1990.

Nelson, Gary G., 1992. Adaptive, Multi-Scaled ATM: Making it Work, USDOT, FAA, ARD-100.

Nelson, Gary G. and Peter M. Allen, Self-Organizing Geography: Scaled Objects and Regional Planning in the U.S., NECSI Bi-Annual Meeting, 2002, Nashua, N.H.

Nelson, Gary G., 2003. "Next TFM: Collaboration to 2015 and Beyond", for the Federal Aviation Administration (FAA), AUA TAC.

Nelson, Gary G., 2005. High-Level Architecture of Homeland Security, Homeland Security Institute.

Nelson, Gary G., 2006a. Organizational Evolution, Life-Cycle Program Design: Essential Issues in Systems Engineering and Acquisition of Complex Systems. Presented at Capital Science 06 (Arlington, VA) and to be published in the Proceedings of the Washington Academy of Sciences.

Nelson, Gary G., 2006b. A Risk-Managing Enterprise as a Complex Adaptive System. Presented at the Risk Symposium 2006, Santa Fe, NM.

Norman, Douglas O. and Michael L. Kuras, 2004, Engineering Complex Systems, The MITRE Corporation, January 2004. At http://www.mitre.org/work/tech_papers/tech_papers_04/norman_engineering/

Rittel, H., and M. Webber, 1973; "Dilemmas in a General Theory of Planning" pp 155-169, Policy Sciences, Vol. 4, Elsevier Scientific Publishing Company, Inc., Amsterdam. Wikipedia summary at http://en.wikipedia.org/wiki/Wicked_problems

Rosenkrantz, Daniel J. et. al., 2004, Structure-Based Resilience Metrics for Service-Oriented Networks.

RTCA 2002. National Airspace System Concept of Operations and Vision for the Future, Washington, DC.

Salthe, S.N., 1985. Evolving Hierarchical Systems: Their Structure and Representation, Columbia University Press.

Salthe, S.N., 1993. Development and Evolution: Complexity and Change in Biology, MIT Press.

Salthe, S.N., Summary of the Principles of Hierarchy Theory, November 2001. At http://www.nbi.dk/~natphil/salthe/hierarchy_th.html

Simon, H.A. and A. Ando, 1961. Aggregation of variables in dynamic systems. Econometrica: 29:111−138, 1961.

Tumer, K., and Wolpert, D.H., 2000. “Collective Intelligence and Braess’ Paradox”, in Proceedings of AAAI 2000, Morgan Kauffman. Available at http://ic.arc.nasa.gov/people/dhw/collectives.php

Wardrop, J. G., 1952. Some theoretical aspects of road traffic research, Proceedings, Institution of Civil Engineers, PART II, Vol.1, pp. 325-378.

Wolpert, D.H., K. R. Wheeler and K. Tumer, Collective Intelligence for Control of Distributed Dynamical Systems, NASA Ames Research Center, Tech Report: NASA-ARC-IC-99-44 August 16, 1999. Available at http://ic.arc.nasa.gov/projects/COIN/Pubs/bar_europhys2000.pdf See also the collection of Wolpert papers at http://ic.arc.nasa.gov/people/dhw/collectives.php