Asymmetric Threat Assessment Tool (ATAT)

Gary W. King, Matthew Schmill, Andrew Hannon. Paul Cohen
140 Governors Drive
University of Massachusetts, Amherst
Amherst, MA 01003
413 577-0176
gwking@metabang.com
{schmill, hannon}@cs.umass.edu
cohen@isi.edu

Keywords: Knowledge acquisition/engineering, synthetic environments, group and individual behavior, intelligent agents

ABSTRACT: *Planning for Effects Based Operations (EBO) requires understanding not just immediate effects – bomb damage and casualties – but also non-immediate secondary effects. This is difficult because secondary effects involve complex interactions of political, military, economic, social, infrastructure and information (PMESII) factors. Indeed, our experience in Bosnia, Somalia and Iraq implies that the psychological hearts-and-minds effects are often of greater importance in the long run than the primary military ones. We present the Asymmetric Threat Assessment Tool (ATAT), an agent based society simulator designed to explicitly model interacting causal chains on both physical and psychological levels. We discuss the use of a prototype of the tool in a Joint Forces Integrated Battlefield Command (IBC) experiment where it provided "what-if" analysis of situation dependent actions and reactions.*

1. Introduction

Planning for Effects Based Operations (EBO) requires understanding not just immediate effects – bomb damage and casualties – but also the non-immediate secondary ones. This is difficult because secondary effects involve complex interactions of political, military, economic, social, infrastructure and information (PMESII) factors. Indeed, our experience in Bosnia, Somalia and Iraq implies that the psychological hearts-and-minds effects are often of greater importance in the long run than the primary military ones. To model and predict PMESII effects, we have developed an agent-based society simulator, which incorporates both physical and psychological factors.

The Asymmetric Threat Assessment Tool (ATAT) was developed to help planners assess the secondary effects of both military and nonmilitary actions. These actions are the inputs to ATAT, which propagates their PMESII effects through a simulated society. ATAT provides time-series graphical representations of these effects.

1.1 The Simulated World

ATAT consists of a 2-dimensional world with population centers, abstract factories and agents (see Figure 1Figure 1).

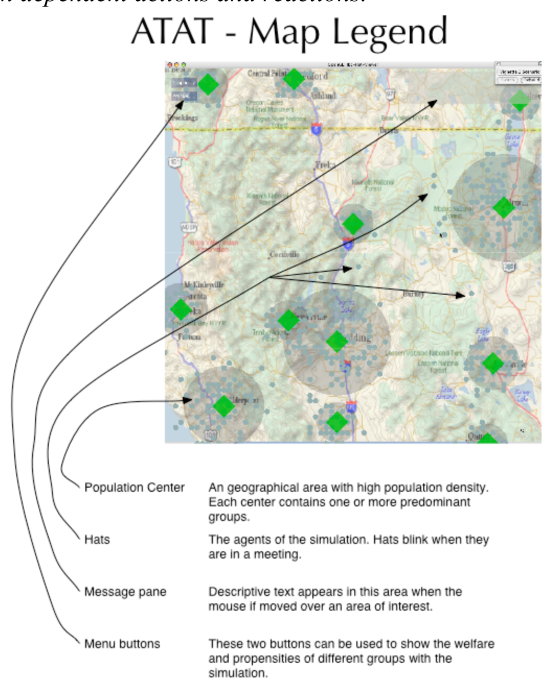

Figure 1: Main features of the ATAT map viewer

Population centers and factories are connected by links over which materials move[1]. Each agent has a home in a population center. Factories consume and produce materials and resources that move over links. Agents work at factories.

[1] Materials and supply are not explicitly modeled in the ATAT prototype.

Agents may also belong to numerous groups, of which there are many types: social, family, ethnic, religious and so on. Agents have dispositions towards other agents based on their group memberships. Agents also have propensities towards action and attributes describing their state (happiness, welfare, and so on). Dispositions, propensities and attribute values all change as the simulator runs.

1.2 Events, Reactions, Meetings, Opportunities, Actions

During a simulation run, the player takes actions such as degrading a factory, attacking a militia group, or building a new population center. These actions produce events in the simulation. Agents who directly experience events have strong

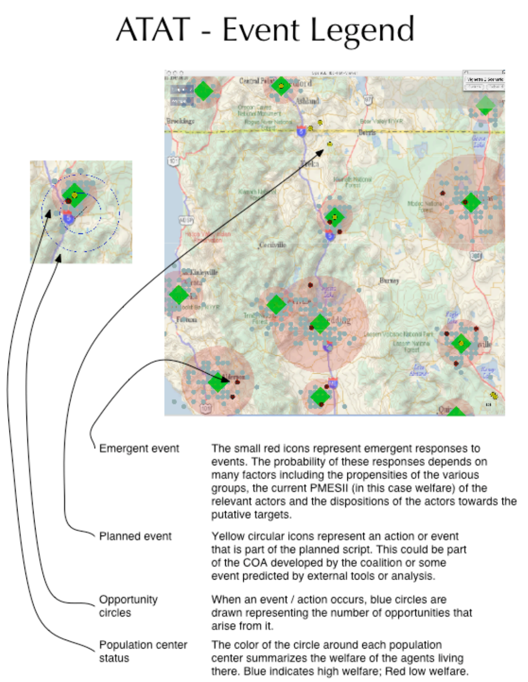

Figure 2~~2~~: Events and opportunities in ATAT

reactions to them. Agents who experience events indirectly have weaker reactions. An agent's reaction depends on many factors: the kind of event, the perpetrator, the agent's disposition towards the perpetrator and so forth. All reactions perturb an agent's dispositions, propensities and attributes (DPA). These perturbations are modeled by a simple set-point mechanism (see section 2 below).

Agents meet continually, and at meetings they influence each other's DPA. These influences depend on the agent's dispositions towards one another and on other agent attributes like welfare, leadership, and initiative[2]. Meetings between members of a group ensure that the DPAs of the group's members are somewhat entrained.

Events also produce opportunities (see Figure 2~~Figure 2~~). For example, a troop movement event produces the opportunity to attack the troops as they move. Opportunities provide groups with the *chance* to act. Groups will act if their propensities towards action are high enough. A group's action may produce another event that will then produce more opportunities and possibly more reactions. This results in a cycle of activity in the simulation as shown in Figure 3~~Figure 3~~.

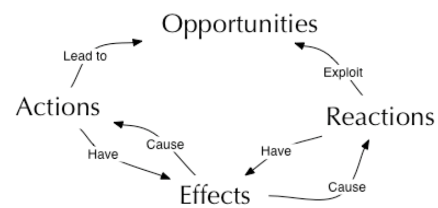

Figure 3~~3~~: the action / opportunity / reaction cycle

In summary:

- Agents witness events directly or hear about them indirectly
- Witnessing events leads to changes in agent DPA.
- DPA is also changed during meetings between agents
- Events produce opportunities for action
- If a group's propensity for an appropriate action is high enough, it will act, producing another event.

A QuickTime movie of the ATAT in action is available on the web (see eksl.cs.umass.edu/~hannon/ibc-demo.mov). The movie shows the simulation running and demonstrates the user interface.

2. Implementation details

2.1 Modeling entities

Population centers, groups, resources, and agents are modeled as objects in the Common Lisp Object System (CLOS). We separate the static and dynamic information for each entity into two classes. The static information describes the entity and provides the information required to build one or more copies of it. Dynamic information includes simulation and time

[2] The ATAT provides a mechanism to implement many different psychological models for agents. The one currently implemented and described here is quite simple.

dependent attributes. For example, a population-center definition specifies the number of agents living in the center, the groups of these agents, their relative sizes and so on. The population-center entity, on the other hand, refers to the actual agents living in the center for the current simulation.

Population centers are connected in a directed graph. Resources flow between centers along these edges (possibly at different rates).[3] Damage to a population center or to the edges alters the flow rate; loss of resources reduces the ability of a population center to support its population and to generate its own resources. Thus, damage in one area propagates out along the graph, placing strain not only on the local population, but on the capacity of surrounding communities to support the wellbeing of their own populations (cf. the failure of the US and CAN power grid, August 2003).

Groups are also modeled as graphs where the vertexes are individuals and the edges represent connectivity. These graphs are meant to be dynamic with both the edges and the vertexes changing over time as agents alter their dispositions towards one another and their membership in different groups. Note that groups can represent any kind of social relation including families, friendships, work, religious and ethnic ties. Each of these has different characteristics in terms of information flow, dispositions and influence. In some societies, religious links are dominant whereas in others, it may be those of family or friendship.

~~2.~~2.2 The Set-point model

The set point model assumes each agent has one or more attributes (e.g., sense of well-being) and each attribute has a set point $\boldsymbol{s}$ and a current value $\boldsymbol{m}$. Events perturb $\boldsymbol{m}$ and $\boldsymbol{s}$ differently. The intuition is that $\boldsymbol{m}$ can be bounced around by events but it will always drift back toward $\boldsymbol{s}$, while $\boldsymbol{s}$ itself shifts in response to events much more slowly. We will focus more on changes to $\boldsymbol{m}$ and $\boldsymbol{s}$ than on their absolute values. The effects of events on $\boldsymbol{m}$ and $\boldsymbol{s}$ depend on the sign (positive or negative) and the magnitude of the events, or rather, on perceived signs and magnitudes. We let $\boldsymbol{p}_t$ represent these quantities at each time t. Because the model is discrete-time, there might be no event in a given time interval $\boldsymbol{t}$. In this case $\boldsymbol{p}_t = 0$. The two main equations governing $\boldsymbol{m}$ and $\boldsymbol{s}$ are:

(1) $m_{t+1} = m_t + \alpha_m p_t + \beta\, sign(m - s)$
(2) $s_{t+1} = s_t + \alpha_s p_t$

The first equation shows that the change in $\boldsymbol{m}$ has two components: one is $\boldsymbol{p_t}$ scaled by the constant $\boldsymbol{\alpha_m}$ and the other is a "drift" toward $\boldsymbol{s}$ scaled by the constant $\boldsymbol{\beta}$. The second equation shows that $\boldsymbol{s}$ changes by a fraction of $\boldsymbol{p_t}$. This fraction is the constant $\boldsymbol{\alpha_s}$ and it should be small; that is, $\boldsymbol{\alpha_m} \geq \boldsymbol{\alpha_s}$. The model ensures that attributes change at two different time scales: $\boldsymbol{m}$ represents daily highs and lows; $\boldsymbol{s}$ represents trends. More complex models are obviously possible but even this simple one exhibits complex and interesting behavior.

3. Inputs and Outputs

ATAT requires a model of population centers, factories and agents. This model includes geographic data, dispositions and propensities toward action, types of actions, types of events, how events produce opportunities and so forth.

A particular simulation run optionally requires a set of exogenous events (e.g., a military course of action, or COA). These events help to drive changes in the simulation. As the simulation runs, it collects time series of DPA by agent. group and geographic area. It also collects time series of events, opportunities and reactions. These data let us explore causal chains of events and better understand how effects move through the agent population.

The simulation can also answer "what if" questions such as: which (re)actions are likely if such-and-such an event occurs. E.g., "What might happen if I move troops through this town?" Monte Carlo simulation can be used to provide probabilities and confidence intervals of reactions given particular COAs.

4. Metrics and Goals

The results produced by the ATAT depend largely on the models and knowledge engineered into it. The ATAT provides a framework for expressing details of the simulation (events, reactions, agent attributes) that must be populated by knowledge about a domain and with models of how agents in the domain react. The prototype described here contains a subset of useful events, actions, attributes and features that we plan to extend.

Metric	Current	Planned
Number of agents	1000	100,000
Number of resources	100	10,000

[3] Instead of modeling materials and supply, the ATAT prototype models population-center efficient as a single variable that is dependent on the efficiency of the population-centers to which it is connected. Though simpler, this model retains the important characteristics of an interconnected society.

Metric	Current	Planned
Number of events (in one simulation run)	50	500
Simulation Speed		
Number of event types	20	100
Number of resource types	1	50
Number of agent behaviors	2	50
Number of group behaviors	20	50
Number of agent attributes	1	10
Number of group attributes	1	10

Table 1: ATAT Metrics and Goals

Table 1~~Table 1~~ shows the current capabilities of the ATAT and our plans for the future. Adding additional event types, group and agent behaviors and attributes will improve the verisimilitude of the simulator and help to model the entire range of COA actions and PMESII effects.

5. The ATAT in use: JFCOM Experiment

5.1 Experiment Overview

The ATAT was one of several tools used in the Integrated Battle Command (IBC) evaluation experiment at Joint Forces Command in Norfolk, VA (JFCOM) on November 15th through the 18th, 2004. The goal of this experiment was to demonstrate how a wide range of decision support tools could be flexibly combined into a single confederation. Communication between tools was managed via the Structured Data Web. This combines a lightweight IPC mechanism with persistence and browsable semantic metadata (Dyer, 2004 a; Dyer, 2004 b).

The experiment methodology was to have two groups of military planners prepare situation assessments and Course of Action (COA) evaluations. One group would use the regular Military Decision Making Process (MDMP) whereas the other would have all of the tools at their disposal. Tool developers operated their own products so that the military planners could focus on their usual tasks without having to worry about learning new technologies.

A detailed scenario was prepared for the JFCOM experiment. This involved the US Military in peacekeeping operations between two fictitious nations. The scenario included both traditional force-on-force actions and asymmetric threats by terrorists and insurgent elements. Each day of the experiment began with a presentation of the scenario's daily event summaries (InSums). The planners then assessed the situation and briefed their commander. The commander then decided how he wanted to move forward and presented the planners with his guidance. Finally, the planners developed and evaluated COAs and presented these to the commander.

5.2 ATAT's Role: the plan

The ATAT was designed to assist in COA planning and evaluation. We used data similar to Operational Net Assessment[4] to populate the ATAT with information about the towns and cities, groups, their dispositions towards one another and their propensities for action. We also used ONA-like data to prepare for the kinds of actions and opportunities we expected to see within the COAs (e.g., fire fights, demonstrations, roadside bombings and the like). Finally, we scripted the relevant insurgent and coalition events from the daily InSums. Our plan was to simulate COAs in order to produce qualitative predictions of the PMESII effects they would be likely to produce.

5.3 ATAT's Role: what really happened

Because the IBC experiment was part of a rapidly evolving rolling start, the scenario continued to change right up to the experiment itself. This meant that some details were added to the scenario that we were unable to effectively simulate.[5] We were, however, still able to model the COAs presented with some degree of fidelity and to use these models to make predications regarding PMESII effects.

An advantage of simulation is that we can make explicit our assumptions. For example, Figures 4-6 show screenshots of the ATAT running with events from the IBC InSums. They demonstrate how the ATAT was tuned to model different assumptions. In these figures,

---

[4] Operational Net Assessment (ONA) - A continuously updated operational support tool that provides a Joint Task Force commander visibility of effects-to-task linkages based on a "system-of-systems" analysis of a potential adversary's political, military, economic, social, infrastructure, and information (PMESII) war-making capabilities.

[5] For example, the scenario unexpectedly gained two additional nations, one of which supported the US role and one of which was highly critical. The ATAT prototype had been designed to model the effects in populations of single countries, not as a simulator of geopolitical relations.

yellow (light gray) circles show events scripted from the daily InSums whereas red (dark gray) circles (see especially Figure 5) show events that emerge from the interactions of the simulation. They show areas where agents with propensities for insurgent behaviors are known to act and have significant probabilities of doing so.

Together, the three figures show the simulation under optimistic, pessimistic, and expected levels of agent volatility. The many red reactive events in Figure 5 show how a more volatile population produces more insurgent activity.

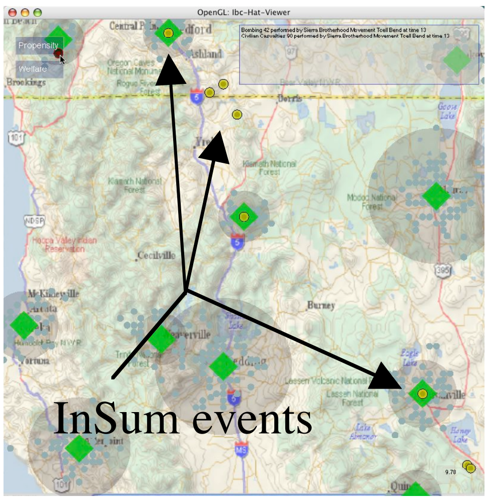

Figure 4: best case assumptions

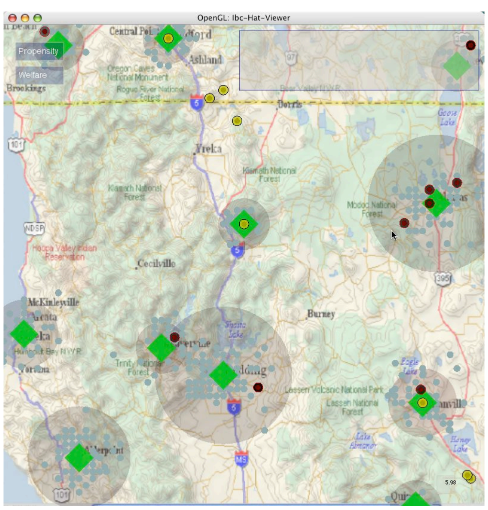

Figure 5: Worst case assumptions

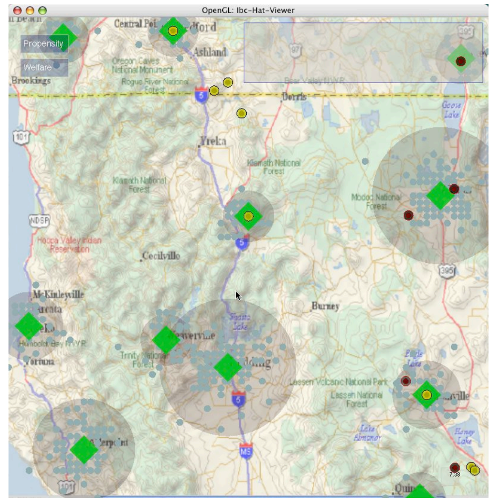

Figure 6: Expected agent volatility

5.4 COA Evaluation

To evaluate COAs, we would add events to the existing InSum script for that day. For example, one COA used strong economic stimulation beginning at day 14 to attempt to reduce insurgent activity by winning the “hearts and minds” of the native population. Figure 7 shows the simulation at day fourteen.

The next two figures (Figures 8 and 9) show the simulation at day 25 with and without the proposed COA.

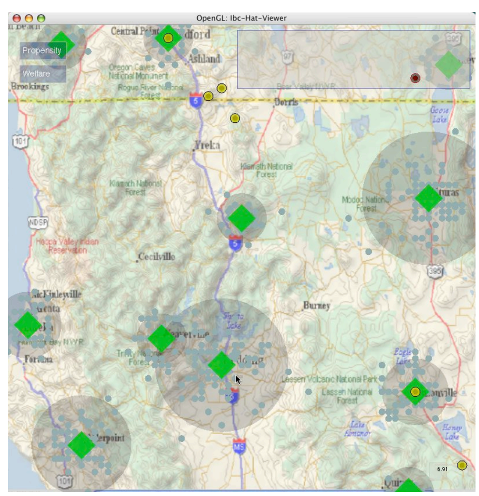

Figure 7: COA Evaluation, Day 14

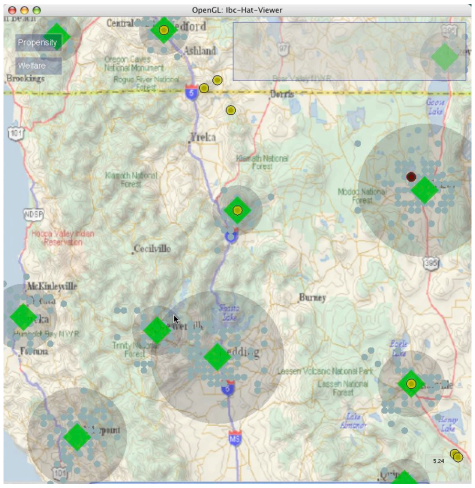

Figure 8: COA Evaluation; Day 25 with economic stimulation

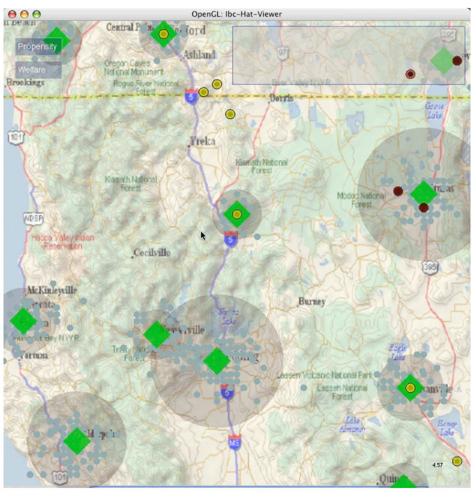

Figure 9: COA Evaluation; Day 25 without economic stimulus

Comparing the two figures shows that the addition of an economic stimulus greatly reduced the probability of insurgent activity. This information was received well by the political-military planner and included in the day's final briefing.

6. Future Work

The IBC experiment indicates that tools like the ATAT can be a useful part of COA evaluation but leaves open many research questions. The architecture of the ATAT is partially logic-based and can support diverse models of actions, opportunities, group dynamics, and psychological activity. Knowledge engineering of such models with the help of Subject Matter Experts (SME) is an open and ongoing research problem (see for example: Clark *et al.* 2001; King *et al.* 2003a, 2003b). Regardless of the models, psychological simulations are very difficult to verify and validate. We can, however, use Monte Carlo analysis to develop confidence intervals on the simulation results. We are also interested in developing dynamical phase portraits of the behavior of attributes, dispositions and propensities over time. Such portraits should both help to validate the simulator as well as summarize behavior in a way that captures its dynamics.

7. Summary

Effects Based Operations require a perspective that includes not only immediate effects but also the non-immediate and often non-linear secondary ones. These secondary effects arise from multiple factors interacting along multiple pathways both within single levels of the PMESII spectrum and across levels. For example, political effects can engender social changes which lead to altered economic activity that may have its own political ramifications. These complex and recurrent causal chains are exceptionally difficult to handle analytically but are relatively easy to model in simulation.

The ATAT approach to modeling effects is effective because it:

- Provides an abstract model of terrorist activity as a dynamic process where coalition and insurgent actions are interpreted subjectively by the agent population.
- Enables "what-if" analysis of situation dependent actions and re-actions
- Models the terrorist adversary as an enemy who takes actions for their direct and indirect effects.
- Aids in training for non-linear effects, unexpected consequences, and cross-spectrum PMESII reasoning.

The prototype implementation of the ATAT used in the IBC experiment included enough models of different groups, actions and reactions to successfully guide planners when making decisions about where, when and how to operate. By using agent-based simulation and tracing chains of causes to effects, the ATAT shows how the probabilities of different outcomes change with coalition activity.

In summary, the ATAT provides a means to analyze how small actions can achieve large ends and how coalition forces can best plan to achieve desired PMESII effects.

8. References

P. Clark, J. Thompson, K. Barker, B. Porter, V. Chaudhri, A. Rodriguez, J. Thomr, Y. Gil, and P. Hayes. (2001). Knowledge entry as the graphical assembly of components: The SHAKEN system

D. Dyer. (2004). On the Structured Data Web as infrastructure for web services. Retrieved from the web May 27, 2004. http://activecomputing.org/papers/Structured-Data-Web.doc

D. Dyer. (2004). Structured Data Web: A web-like method for sharing and interpreting structured information for bottom-up development. Retrieved from the web June 1, 2004. http://activecomputing.org/papers/SDW-as-infrastructure-for-Web-Services.doc

G. W. King, C. T. Morrison, and P. R. Cohen. (2003). Action models. In S. Chick, P. J. Sanchez, D. Ferrin, and D. J. Morrice, editors, Proceedings of the 2003 Winter Simulation Conference.

G. W. King, C. T. Morrison, D. L. Westbrook, and P. R. Cohen. (2003). Bridging the gap: simulations meet knowledge bases. In Enabling Technologies for Simulation Science VII. Edited by Sisti, Alex F.; Trevisani, Dawn A. Proceedings of the SPIE, Volume 5091, pp. 383-394 (2003), pages 383–394, 2003.

Author Biographies

GARY W. KING is a Senior Research Fellow at the University of Massachusetts, Amherst and the lab manager of the Experimental Knowledge Systems Laboratory under the direction of Paul R. Cohen. His research interests include programming language design, planning, simulation, language acquisition and algebraic information theory. His web address is www.metabang.com.

MATT SCHMILL is a postdoctoral research associate at the University of Massachusetts, Amherst. His core area of research is artificial intelligence, specializing in behavior and cognitive development.

ANDREW HANNON is a computer programmer with the Experimental Knowledge Systems Laboratory at the University of Massachusetts, Amherst. Mr. Hannon holds a B.S. in Computer Science from the University of Massachusetts, Amherst. His research interests involve simulation, computer graphics, and computational geometry. He can be reached by email at hannon@cs.umass.edu.

PAUL COHEN directs the Center for Research on Unexpected Events at USC's Information Sciences Institute (ISI). He is a research professor of Computer Science at USC and deputy director of the Intelligent Systems Division at ISI.