SEPO
Volume 6 – Number 1
COLLABORATIONS
Inside Knowledge for MITRE's Systems Engineering Community
www.mitre.org/work/sepo/

Agent-based Modeling

This issue of *Collaborations* explores how researchers at MITRE are creating innovative solutions in the realm of agent-based modeling (ABM) and illustrates the varied approaches being taken across domains, problems, and sponsors.

Agent-based modeling, which describes a system from the viewpoint of its constituent parts, is most useful when trying to understand a complex system by observing how its individual components interact. For instance, ABM is often used by researchers to analyze human systems. While human systems are composed of many, separate, interacting components, the behavior of each part can be broken down into simple terms and simulated using ABM.

A frequently cited example is a traffic jam, in which the individual drivers serve as the agents and their driving behaviors—level of aggression, frequency of changing lanes, whether they pass on the right—are the rules for the ABM model. Using this model, the drivers' behaviors can be observed while different events, such as a forest fire near a highway, are introduced. It is then possible to see how these events affect emergent properties of the traffic jam and then use these observations to predict what might happen in a future traffic jam.

The first set of articles presents different applications of a MITRE-based methodology for designing creative solutions to complex systems engineering and decision support problems. These applications include developing an agile analysis capability for the US Army; using combinations of sensors and personnel to protect large venues from contraband materials; creating a simulations toolkit for counterinsurgency analysis; and modeling the effects of the pandemic flu on telecommuting.

The MITRE Epidemiological Model is then presented, followed by several examples of how it is being used to manage the consequences of a pandemic influenza. One article uses the model to explore how different combinations of factors affect a pandemic flu outbreak aboard Navy vessels. Another discusses how the model can be used to evaluate planned governmental responses to a flu outbreak.

In the last two articles, consideration is given to the question of validation and agent-based modeling, which takes on particular challenges when applied to complex adaptive systems, especially those involving human social interactions. And finally, the DAP-E model is presented, a simulation technique that represents social and attitudinal factors in ABMs, which may allow for more realistic results.

Contents

An Agile Methodology for Addressing Complex Systems Engineering and Decision Support Problems

Matthew Koehler and Brian Tivnan

Designing complex systems to accommodate human behavior is a challenging task. To aid our sponsors faced with this challenge, MITRE has created a methodology—as well as the tools, techniques, and infrastructure to support it—for developing agile solutions to decision support problems. Known as the Infrastructure for Complex Systems Engineering, our methodology and toolkit begin with initially defining the problem, eliciting data from sponsors and/or subject matter experts (SMEs), and gathering additional data from primary and secondary sources. The MITRE team then moves the data and problem statement into a prototyping environment, where a quick-cycling process with SMEs and sponsors helps to flesh out questions and better define how a solution might look. As the prototypes mature, they move into a robust, scalable system that can more easily be run in a high performance computing environment to explore a vast sample of the model's parameters.

Because these scenarios hinge on human behavior, the prototyping and robust systems used are agent-based models—they focus on explicit representations of individual actors within the system, and highlight their interactions with other actors and the environment. At the same time, we can integrate high-resolution physics-based models that may require a high level of verisimilitude. The models can also be broken down into individual parts that can be used separately depending on each customer's needs. The final challenge is to analyze and present the resulting large amount of data so it can be used effectively by different audiences: MITRE, for development and debugging purposes; SMEs, who understand the topic but may not be well-versed in modeling and data analysis; and the decision-maker, who may not be an SME in the model's area nor understand modeling.

DESIGNING COMPLEX SYSTEMS TO ACCOMMODATE HUMAN BEHAVIOR IS A CHALLENGING TASK.

The following four articles explore how this methodology is being used for different purposes with different sponsors.

Creating Responsive Analysis Alternatives for the U.S. Army

Matthew Koehler, Sarah Johnson, Zoe Henscheid, Carl Burke, Khalid ElArini, Neal Rothleder, Adam McLeod, and Lyndon Brown

During budget planning, the G8/FDA (U.S. Army's G8/ Force Development/Warfighting Analysis Office) is asked to respond to many "what if" scenarios: "What if we do not have the doctrinal number of smart mortar rounds?" "What would happen if we had more laser guidance systems?" Not only are these difficult questions to answer, but the G8/FDA is usually expected to do so within 24 to 48 hours.

Traditionally, the G8/FDA would turn to a subject matter expert, whose innate understanding of the situation would be applied to come up with the best possible answer in the limited time available. The G8/FDA was not satisfied with this process, however, and wished to increase its analytic rigor. A number of Army analysis groups were more than willing to help, but could not complete studies in the necessary timeframe. The G8/FDA then asked MITRE for help creating an in-house analysis capability.

Tailoring the Methodology to the G8/FDA

To help create this analysis capability, we applied our agile systems engineering and decision support methodology and infrastructure, using only the pieces relevant for the G8/FDA's needs. For example, because the system needs to be extremely responsive and organic to G8/FDA analysts, we removed the high-resolution agent-based model and the high-performance computing environment. Since the G8/ FDA wanted to augment rather than supplant its existing analytic process (i.e., it is looking for rough order-of-magnitude answers: "Is this better or worse?"), the analysis was stopped at the prototype stage.

Off-the-Shelf Scenarios and Output Analysis Tools

Because an agent-based modeling framework already existed, we concentrated our efforts on developing two other important features: off-the-shelf scenarios and output analysis tools. Off-the-shelf scenarios give G8/FDA analysts a "warm start" for responding to questions more quickly. Some of these include infantry brigade combat team fights, heavy brigade combat team fights, convoy movement, small scale urban combat, and rotary wing aircraft sorties.

We also worked on creating a useful set of output analysis tools so G8/FDA analysts could produce responsive, intelligible results. Our analyses focused on casualties suffered by both sides, categorizations of casualties suffered, and shots fired; we then developed a series of visualization tools.

For example, Figure 1, shows a time series of casualties suffered by Blue (U. S. Army) forces in a set of simulation runs. The figure contains a large amount of information: lines show the minimum, maximum, and mean value of casualties over time, and a jittered scatter plot indicates the number of casualties suffered at each time step in every run. Showing the casualty values along with the minimum, maximum, and mean gives a much more meaningful picture of the actual distribution of casualties, in a way that summary statistics cannot.

Visualization Tools

Another visualization tool we developed is an "agent interaction table," which displays mean and distributional data for interactions that took place between any two agent categories. For example, Figure 2 shows a killer/victim (red/blue) interaction table; only mean values are displayed. An interaction is defined as any encounter between two entities that may affect the state of one or both. Any table entry represents the average number of interactions triggered or instigated by the entity at the left-end of that row with the entity at the top-end of that column. To help meaningfully visualize large numbers of runs, we developed an analysis tool for looking at many user-defined simulation runs all at once, by layering them on top of each other and making each slightly transparent. In this way, an analyst can see the average behavior of the agents as well as outlier behavior.

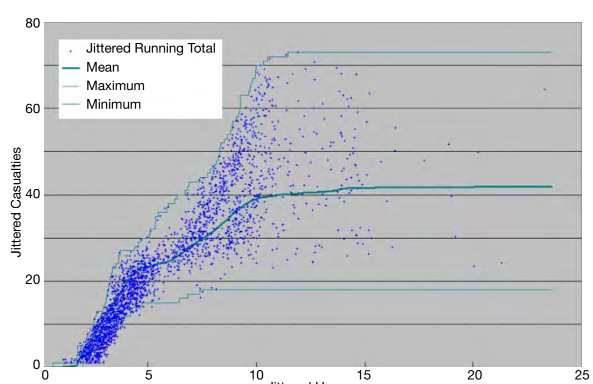

Figure 1. Jittered Casualties over Time

We are also exploring methods for displaying the correlative structures among agent movement, success or failure, and spatial relationships. Our current capability, called delayed outcome reinforcement plotting, combines machine learning with density playback. Essentially, the paths of successful agents are positively reinforced, while the paths of failing agents are negatively reinforced. The values are then assigned a color and displayed as a movie, showing agent flows around terrain colored by their success or failure.

Data Mining

Finally, we have begun to explore some very sophisticated data mining techniques to categorize the state space of the simulation runs into a chain of meaningful events. This will allow us to create analyses that are more immediately useful to decision makers and subject matter experts—for example, being able to explicitly say, "Blue does better when squad 10 provides covering fire for squad 6." At the same time, we are transitioning all of these capabilities to the G8/FDA for their use on in-house analyses by the end of this fiscal year.

	Blue ASLT	Blue Air	Blue Avn	Blue ndirec	Blue ndirec	Blue ndirec	Blue RSTA	RED BTR	Red CTR	Red orwar	Red orwar	Red ndirec	Red Major	Red lortars	Red RSTA
Blue ASLT	0.03	0	0	0	0	0	0	1.6	0.08	5.37	0.68	0.26	13.06	0.6	0
Blue Air	0	0.08	0	0	0	0	0	0.08	0.01	0.08	0	0.01	0.44	0.15	0
Blue Avn	0	0.01	0.01	0	0	0	0	3.16	0	1.99	0.08	0.12	2.12	0.01	0
Blue ndirec	0	0	0	0	0	0	0	0.06	3.36	0	0	0.08	0	0.08	0
Blue ndirec	0	0	0	0	0	0	0	0	0.44	0	0	0.14	0	0	0
Blue ndirec	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Blue RSTA	0.01	0	0	0	0.01	0	0.04	5.82	0.03	8.59	2.04	0.13	9.41	0.46	0
RED BTR	0.56	0.39	6.95	0.56	0.01	25.42	1.07	0.04	0	0	0	0	0	0	0
Red CTR	0	0	0.26	0	0	0.05	0	0	0	0	0	0	0	0	0
Red orwar	3.48	1.7	1.7	10.32	16.66	0.04	16.1	0.1	0	0.88	0.02	0	0	0.01	0
Red orwar	5.57	0.14	0.28	2.19	7.5	0.04	8	0.08	0	0.13	0.01	0.04	0.07	0.02	0
Red ndirec	0	0	0.15	0	0	0.01	0	0	0	0	0	0	0	0	0
Red Major	19.88	2.2	6.29	99.46	3.94	0.12	14.06	0	0	0	0	0	1.02	0	0
Red lortars	0.16	0.1	0.04	0.82	0.16	4.49	0.2	0	0	0	0	0	0	0	0
Red RSTA	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Figure 2. Agent Interaction Table

For more information and related reference materials, contact Matthew Koehler at mkoehler@mitre.org, or 703-983-1214.

Helping the Department of Homeland Security Protect Airports and Passengers

Philip Barry, Matthew Koehler, and Garry Jacyna

Identifying Passengers Carrying Explosives

By the year 2015, Los Angeles International Airport expects more than 78 million passengers a year to pass through its facility, up from 68 million in 2000. These numbers demonstrate how critical it is to identify passengers who may be carrying hazardous materials. Historically, portable explosives have been the cause of 65 percent of attacks on airports, and the majority of the 4,280 fatalities resulting from attacks on aircraft. Aggressive use of technology and innovative tactics are critical to mitigate this very real threat.

To support the Department of Homeland Security, MITRE is using an agent-based model integrated with high fidelity sensor models to investigate the best combination of tactics and technology, using the Infrastructure for Complex Systems Engineering, or ICE, methodology. First, we prototyped the scenario, and then moved the prototypes to a more robust agent-based modeling framework in the high performance computive environment (HPC). When it became difficult to obtain sensor data, we incorporated very high-resolution physics-based models of sensors to compensate for the lack of experimental data. Finally, we used the genetic algorithm from the HPC to search for optimal sensor and tactics combinations to best protect the venue.

The Agents

As shown in Figure 3, three types of agents are used in the model: security, civilians, and bad actors. Bad actors model terrorists whose goal is to carry explosives undetected through the turnstiles. In the model, bad actors must traverse a corridor from left to right in which they are likely to encounter sensors and security guards. A bad actor that is able to cross the turnstile is considered successful.

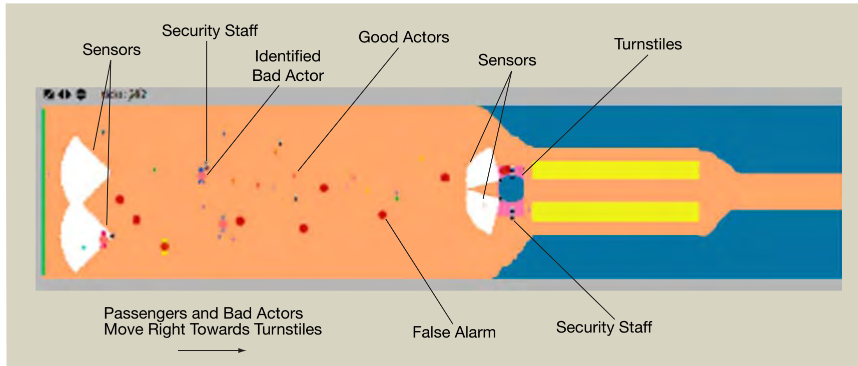

Figure 3. Simulation Set-up

Bad actors will take evasive measures, acting independently, to avoid security guards. For example, this was modeled by including a distance metric that was used to grade the response of a bad actor relative to the proximity of security guards. We have not yet modeled evasive measures where bad actors collaborate among themselves.

Bad actors will also sacrifice themselves if necessary. For example, if security guards close off potential escape routes and the bad actor considers the situation hopeless, he will exhibit sacrificing behavior and detonate the explosives.

Security agents are stationed throughout the corridors and around the turnstiles. Their primary goal is to interdict the bad actors. Security agents use fused data from the sensors to evaluate whether a passenger is a bad actor or not. Since false alarms are possible, civilians may also be targeted and stopped by security agents.

The remainder of the agents are "innocent passengers" who are represented as a rate per unit time. The number created in each time-step is drawn from a random-exponential distribution. A small percentage of passengers (0.005%) may then be further instantiated as bad actors.

The Sensors

The model uses three types of sensors: two passive millimeter wave (MMW), and one infrared. These sensors are placed in fixed locations around the environment to detect and classify concealed weapons under a person's clothing. Given the public nature of the space inside an airport, it is often desirable to screen people from a distance to reduce the likelihood of long lines or unnecessary crowding. To achieve this, there are important problems that remain to be solved regarding the composition, placement, and tasking of semi-autonomous sensors; extracting and using actionable information from these systems; and developing a reasonable concept of operations.

The first stage of MITRE's effort involved developing passive MMW sensor models to accurately predict the probabilities of detecting and classifying concealed weapons at range. In the MMW band, objects are described by their apparent temperature—a combination of the actual temperature and that of the reflected background. For a metallic object, the apparent temperature is the background temperature. The detection probability is determined using the effective contrast temperature and the number of beams covering the object at a fixed false alarm probability.

Going Forward

Our work so far has demonstrated that integrating behavioral and detailed physical models can influence the concept of operations for protecting a large public venue such as an airport. Our exploratory research shows that when using an optimized sensor placement approach combined with appropriate security guard tactics, a significant percentage of bombs can be detected and interdicted. In addition, we have compared our results to published theory and found good agreement.

The work discussed here will help determine the benefits of a specific sensor placement for a real world experiment scheduled for September/October 2008. Towards that end, we will continue to refine our model and iterate it with the subject matter experts to ensure reasonable behavioral representations. At the same time, we will continue our investigations to gain insight into the optimal combination of sensor placement and security guard tactics.

For more information and related reading materials, contact Philip Barry at pbarry@mitre.org, or 703-983-7826.

Counterinsurgency Sandbox for Decision Support

Matthew Koehler, Brian Tivnan, and Thomas Wilk

Counterinsurgency (COIN) operations require a flexible, adaptive force that can compete against asymmetrical and adaptive adversaries. The complex dynamics of COIN operations require that commanders use analytical tools to assess the potentially far-reaching implications of their operations. However, current force-on-force combat models have two inherent limitations: First, because these models are predicated on force exchange ratios and prescribed doctrinal operations, they cannot capture the dynamics of COIN operations, and therefore cannot incorporate the learning and adaptation that are paramount to meaningfully model COIN operations. Second, current models cannot support rapid prototyping and response because they typically require several months of preparation to execute a single scenario. In response, a MITRE team created the COIN Sandbox—a computational model designed to represent the complex dynamics of COIN operations and to facilitate rapid analysis.

Agent-Based Modeling Approach

An agent-based model (ABM) framework is a natural approach for representing learning and adaptation in complex dynamic systems. The ability to model this is crucial since it is a key aspect of the insurgency the U.S. faces in countries like Iraq. Because we face a highly adaptive adversary, many solutions likely will become obsolete very quickly. An ABM, however, is ideally suited to reflect the dynamics of COIN operations because it can explicitly depict adaptive agents and represent individuals within the operation. Each agent can be assigned appropriate characteristics for behavior, attitudes, role, and social interactions.

> **THIS IS A GAME OF WITS AND WILL.**
> **YOU'VE GOT TO BE LEARNING AND ADAPTING CONSTANTLY TO SURVIVE**
> *General Peter J. Schoomaker, USA, 2004.*

Recently, researchers at the Santa Fe Institute, Argonne National Lab, and other labs have made advances in modeling learning and adaptive phenomena previously unapproachable in complex systems (e.g., financial market

Continued on page 6

bubbles and policing actions). MITRE has been awarded a two-year fellowship with the Santa Fe Institute to conduct collaborative research in support of our sponsors, which will give us direct access to many experts in the field of modeling complex systems.

Rapid Prototypes

MITRE has also successfully developed rapid prototypes of combat and non-combat models for the Army G-8/ Force Developing/Warfighting Analysis office (FDA), the 1st Marine Expeditionary Force (I MEF), and the Marine Corps Warfighting Laboratory. Our work includes modeling changes in civilian perceptions of quality of life and coalition activities during stability and support operations; convoy protection planning; and infantry brigade and heavy brigade combat team scenarios.

Several of these prototypes are appropriate for integration into the COIN Sandbox. For example, Figure 4 illustrates some results from the rapid prototype for I MEF, which explores how information propagation affects attitudes of an indigenous population towards coalition forces. The AgentTracker chart shows the perceived, though highly abstracted, quality of life (i.e., good or bad experiences) for five randomly selected agents. During the run, the agents' perceptions change a great deal, due to both positive and negative events that occur and to agents interacting and "discussing their feelings" about quality of life. The likelihood of agents entering into discussions is a function of

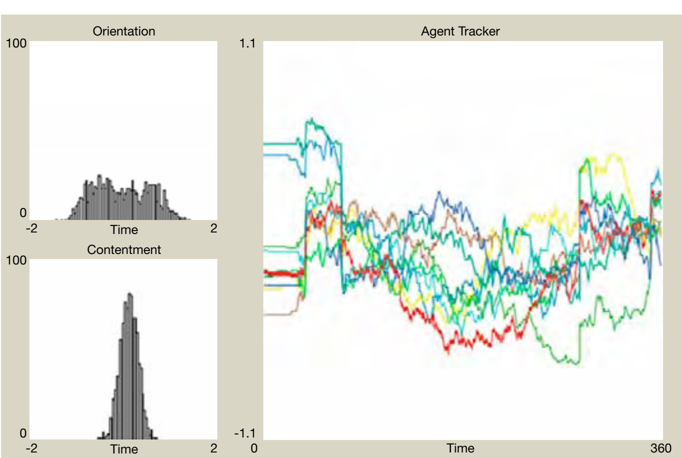

Figure 4. Typical Results from a Run of the I MEF Civilian Sentiment Model

demographic characteristics such as religion, clan, sex, age, and marital status.

The left portion of Figure 4 shows the distribution of agent contentment (perceived quality of life) and orientation (feelings toward the coalition). The perceived quality of life has a relatively tight distribution, but orientation is much more diffuse because it is based on the agents' endogenous,

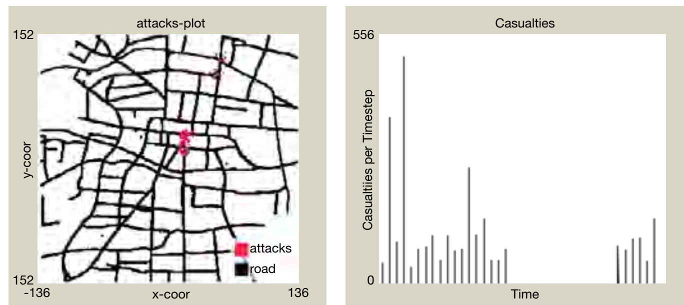

Figure 5. Typical Results from a Run of the Sociocultural Geospatial IED Use Model

heterogeneous feelings about the coalition's legitimacy. If they feel the coalition is illegitimate, their orientations will decline, even when their contentment is improving, because they do not wish to see the coalition succeed.

We also developed a prototype of the sociocultural aspects of improvised explosive device (IED) use and dynamics, which includes weak learning by IED users. As shown on the left in Figure 5, most attacks take place near intersections where the convoy must change direction. An ABM was used to develop this prototype to create a dynamic system in which the bombers chose locations by keeping track of the Blue convoy and trying to find an area where they could best attack. Bombers needed to find roads within neighborhoods that support insurgent activities, while at the same time avoiding roving Blue patrols.

The right side of Figure 5 shows casualties caused by IED attacks. A distribution of casualties was generated that was qualitatively analogous to actual data. Generating relational equivalence to actual data allowed us to say that the model is "on track" and ensure that the primary features and dynamics contained within the model mapped to the primary features and dynamics of the "real" world.

For more information and related reading materials, contact Brian Tivnan at btivnan@mitre.org, or 703-983-3829.

Secondary Effects of Pandemic Influenza: Impact of Telecommuting on Communications Networks

David Bauer, Emmet Beeker, Justin Darkoch, Ernest Page, and Brian Tivnan

Telecommuting is a growing trend, particularly in the services sector of the economy, spurred in part by the adoption of broadband Internet access. A 2006 survey found that approximately eight percent of U.S. workers have an employer who allows them to telecommute full-time, and roughly 20 percent of the workforce engages in telecommuting at least once a month. Despite this growing trend, the feasibility of widespread telecommuting during a pandemic scenario, given the existing infrastructure, has not yet been established.

Several national pandemic plans identify telecommuting as a key component of the national response to pandemic influenza. Telecommuting is recommended as a social distancing mechanism to limit disease spread and allow businesses to continue to function. The Centers for Disease Control advises businesses to consider following telecommuting programs for moderate pandemics, and recommends telecommuting programs as an important component of the planning strategy for severe pandemics.

More specifically, the Homeland Security Council (HSC) has advised businesses to plan for up to 40 percent of their employees to be absent during the two-week peak of a six-to-eight week pandemic wave. However, many HSC stakeholders are concerned that the telecommuting strategy is not technologically feasible with the current telecommuting infrastructure. The interdependent aspects of pandemic response provide an excellent example of a complex systems engineering problem with nested, unintended consequences. Agent-based models provide a useful analytical platform to explore the propagation of a pandemic through a population and the behaviors of network users.

TELECOMMUTING IS RECOMMENDED AS A SOCIAL DISTANCING MECHANISM TO LIMIT DISEASE SPREAD AND ALLOW BUSINESSES TO CONTINUE TO FUNCTION.

Pandemic Telecommuting Strategy

MITRE recently undertook a study to answer two key questions: Will the telecommuting strategy work during pandemic influenza? What preparations can be done to better prepare for telecommuting during pandemic influenza? Taking the general telecommuting strategy advocated in national pandemic plans as a given, we focused our efforts on identifying the technical feasibility of the pandemic telecommuting strategy, and on developing steps to better prepare the nation to support the strategy.

The dynamics between pandemic spread and Internet user behavior are assumed to occur primarily at a metropolitan level. The pandemic is anticipated to spread across the country as a result of people traveling between major cities. Within metropolitan areas, population densities and mixing among populations are expected to largely drive the pandemic spread. From the perspective of telecommuters, the majority of communications traffic likely will be between the telecommuter and the enterprise network, and therefore will be mostly confined within the metropolitan area. Thus, the intrinsic scale of understanding the relationships among pandemic spread, Internet user behavior, and Internet performance is assumed to be at a major metropolitan level.

Modeling Framework

We leveraged an earlier large-scale model of a deregulated electricity market to develop a modeling framework to better understand the relationships among pandemic spread, Internet user behavior, and Internet performance. Our modeling framework included the following: a metropolitan-scale pandemic model based largely on seminal epidemiological research; a network-user model to depict various Internet user profiles; a metropolitan-scale model of the Internet; and a visualization tool based on Google Earth. Depicted in Figure 6, these three models were all instantiated in a large-scale, discrete event simulation platform to take advantage of Command and Control's high performance computer—the HIVE.

The pandemic model used U.S. Census Bureau and Department of Transportation data from the Chicago area to examine the potential movement of the Internet user

Continued on page 8

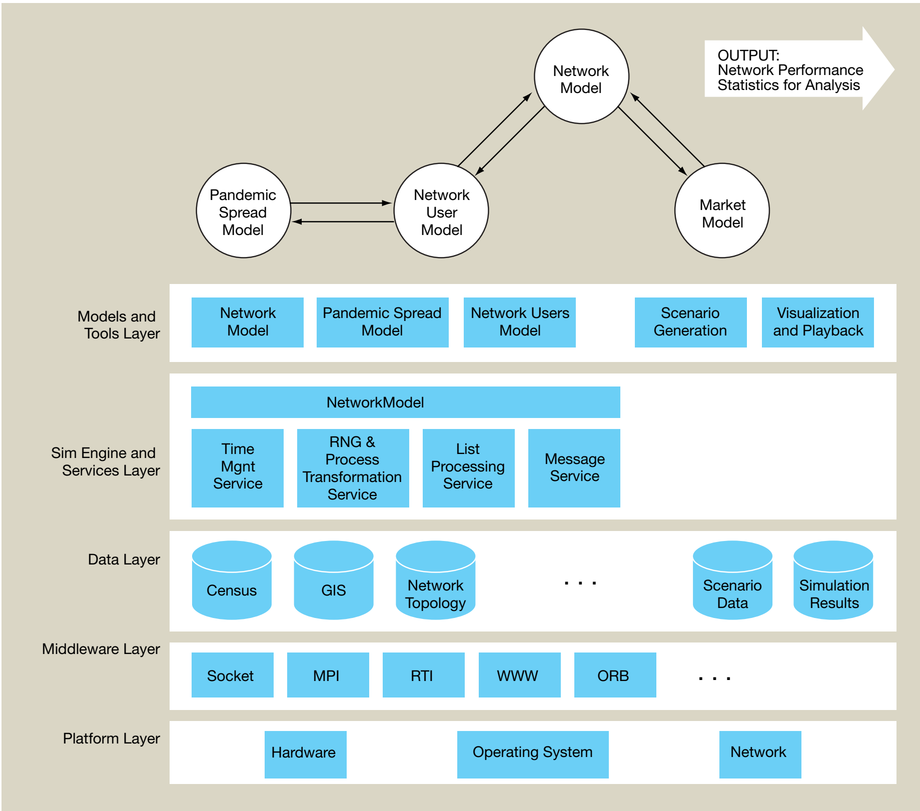

Figure 6. Modeling Architecture for DHS Study

population among home, office, and school locations during a pandemic. Using actual trace-route data, the network model depicted the impact of the change in Internet user behavior on residential Internet access networks. In particular, we examined the ability of telecommuters to perform their functions from residential Internet access networks when increased numbers of users were online during a pandemic. The modeling framework was also used to evaluate the potential impact of phased preparations (e.g., load balancing via shift work, limiting Internet traffic to essential business applications, etc.) to improve network performance.

Defining the scope of the modeling architecture is important for interpreting the model's results. In terms of people and network traffic sources, the models were designed to be scaled to a major metropolitan area. Although the models used data from the Chicago area, the results can be generalized to other metropolitan areas of comparable size.

Since detailed network architectures (e.g., engineering designs for network topologies) are tightly held by network service providers, we used open source information and engineering best practices to construct a typical network architecture model, focusing on representative residential Internet access networks. This modeling framework provided the analytical foundation for a comprehensive study for the sponsor. Through a collaborative research project with the Santa Fe Institute, we are extending the modeling framework by integrating a leading financial market model to analyze the financial impacts of pandemic influenza.

For more information and related reading materials, contact Brian Tivnan at btivnan@mitre.org, 703-983-3829.

The MITRE Epidemiological Model

Emmet Beeker

In 2004, as part of an Innovation Grant, MITRE began using agent-based modeling (ABM) to research disease spread. Promising results led to an Internal Research and Development project in 2005-2006 and the development of the MITRE Epidemiological Model.

As part of our research, we compared and benchmarked our agent-based model against traditional, ordinary differential equation (ODE) models. In the differential equation model, at each time step a portion of the susceptible population became exposed. In the agent-based model, each agent had a probability of catching the disease. When the appropriate probability was used, the results from the agent-based model closely matched those of the ODE model. The ABM was calibrated against widely accepted ODE results, but leading research (e.g., Models of Infectious Disease Agent Study [MIDAS]) has adopted ABM for its ability to represent geo-spatial as well as temporal dynamics.

Using the MITRE Epidemiological (Epi) Model, developed in Repast software, we simulated 800 agents using Susceptible, Exposed, Infected, and Removed (SEIR) disease states. In a SEIR model, the population is susceptible to catching a disease. Then, a segment of the population is exposed to the disease. This exposed population contains the disease in an incubation form, and until these individuals are infected, they can't infect the rest of the population. Eventually, the infected population either recovers or dies, depending on the lethality of the disease. In either case, these individuals are removed from the rest of the population, so that they can no longer play a role in spreading the disease.

Fixed-time Model

This fixed time step model did not scale well—we estimated a one million agent model would take 180 days to simulate. To decrease the simulation time, we scrutinized both the agent-based modeling code and the selected software. We determined the best way forward would be to simulate the disease spread using a discrete event model written in C programming language where we could optimize algorithms and eventually use parallel computing.

Social Network-based

We then changed the time-stepped model described above to a social network-based interpretation using discrete events. Agents moved between locations (e.g., home and work), and each agent's chance of becoming infected was pro portional to the time spent and number of infected people in a particular location.

Instead of checking every agent for infectivity at every time step, agents were checked whenever the group composition changed in a particular location to determine if they were exposed during the preceding time period. This adaptation decreased the simulation time for one million agents from 180 days to 90 minutes and for 10 million agents to 10 hours. Future work includes simulating 300 million agents (i.e., the population of the U.S.) using parallel computing.

MITRE EPIDEMIOLOGICAL (EPI) MODEL

Transition from one disease phase to the next is calculated at a fixed interval using disease equations. In our model, at each time step:

- One agent was removed, and one was added (i.e., simulated birth of a susceptible agent)
- Each susceptible agent selected another agent at random—if the selected agent was infected, then the susceptible agent became exposed (i.e., probability of exposure = number of infected agents/800)
- Each exposed agent had a 0.1 probability of becoming infected
- Each infected agent had a 0.1 probability of recovering.

While the traditional SEIR model uses four disease phases, the MITRE Epidemiological Model can include additional phases, including healthy, exposed, incubating, infected, inoculated, immune, and deceased. MITRE's model is unique in its ability to handle multiple diseases in the population simultaneously and to model heterogeneous mixing of the population. Furthermore, the model can simulate large-scale populations (currently 10 million) and can deliver high execution time performance.

For more information and related reading materials, contact Emmet Beeker at ebeeker@mitre.org, 703-983-3113.

Modeling Pandemic Influenza Virus Outbreaks Aboard Navy Vessels

Janice Ballo, Lynn Cooper, James Diggans, Steven Fairchild, Maeve Kluchnik, Elainne Lea, Meredith Keybl, Mojdeh Mohtashemi, Matthew Peterson, Felicia Sutton, Chrissy Vu, David Walburger, Megan Ward

Controlling infectious diseases in military populations is essential for mission readiness. The recent spread of the avian H5N1 influenza virus and the growing number of human cases have focused attention on influenza pandemic planning efforts. The uncertainties of just where, when, or how a new flu strain will enter human populations require planning efforts that take into account the full spectrum of disease scenarios and medical countermeasures. To facilitate this planning, models must have flexibility to allow end users to "run the numbers" to explore how different combinations of factors will affect the course of an outbreak. The purpose of MITRE's study was to use agent-based modeling to develop and apply such an approach to the population of crew aboard large amphibious assault ships.

Modeling Approach

We used data from the 1918, 1957, and 1968 influenza pandemics to bracket the potential range of disease severity. Traditional SEIR models (i.e., Susceptible, Exposed, Infected, and Removed disease states) allow the end user to manipulate input parameters, such as the number of susceptible people aboard, but they also assume homogeneous mixing of the shipboard population. In other words, every individual on the vessel has an equal chance of coming into contact with an infected person and becoming infected themselves. We developed the Contained Area Stochastic Epidemiological Simulator (CASES) to model the spread of infectious disease aboard naval vessels by adapting the MITRE Epidemiological Model (see page 9). CASES combines a stochastic SEIR-like disease model with an agent-based modeling approach, and can model any number of crew members across any number of ships.

> **THE UNCERTAINTIES OF JUST WHERE, WHEN, OR HOW A NEW FLU STRAIN WILL ENTER HUMAN POPULATIONS REQUIRE PLANNING EFFORTS THAT TAKE INTO ACCOUNT THE FULL SPECTRUM OF DISEASE SCENARIOS AND MEDICAL COUNTERMEASURES.**

CASES Key Features

The CASES model, which was customized for naval applications, can calculate the effects of various countermeasures, including antivirals, protective clothing, patient isolation, and quarantine in single and shared locations. Modeling countermeasure impact on disease spread in terms of timing and percentage of sick crew enables pre-event medical response planning to take place.

CASES also expanded the agent-based features from tracking people between two locations (home and work), to tracking them among any number of shipboard locations (e.g., gym, mess hall, recreational room, etc.). CASES accounts for variations in human behavior by allowing for random movements among locations. We modified the input file format to be human readable, allowing non-technical users to make modifications and assign values to agents' profiles—for example, prioritizing countermeasures for key individuals such as commanding officers. We developed new code for creating schedules and a Java graphical user interface.

CASES allows the end user to manipulate three important factors that contribute to disease spread: crew movement, pathogen characteristics, and countermeasure effectiveness. The crew's movement is modeled using 24-hour schedules. Individual schedules can be customized and repeated on a daily, weekly, or monthly basis. These also permit simulation of random crew member movement.

CASES can also account for numerous pathogen parameters, and can set initial infections either randomly or through a user-specified input file. Finally, CASES can model vaccinations, antiviral medications, protective clothing, and quarantine procedures. The various medical countermeasures can each be assigned effectiveness values, and can be applied to different crew members. For example, certain subpopulations may have been vaccinated or have access to protective clothing, which can be captured in the model. The model also allows individuals to be quarantined in different locations and with variable probabilities.

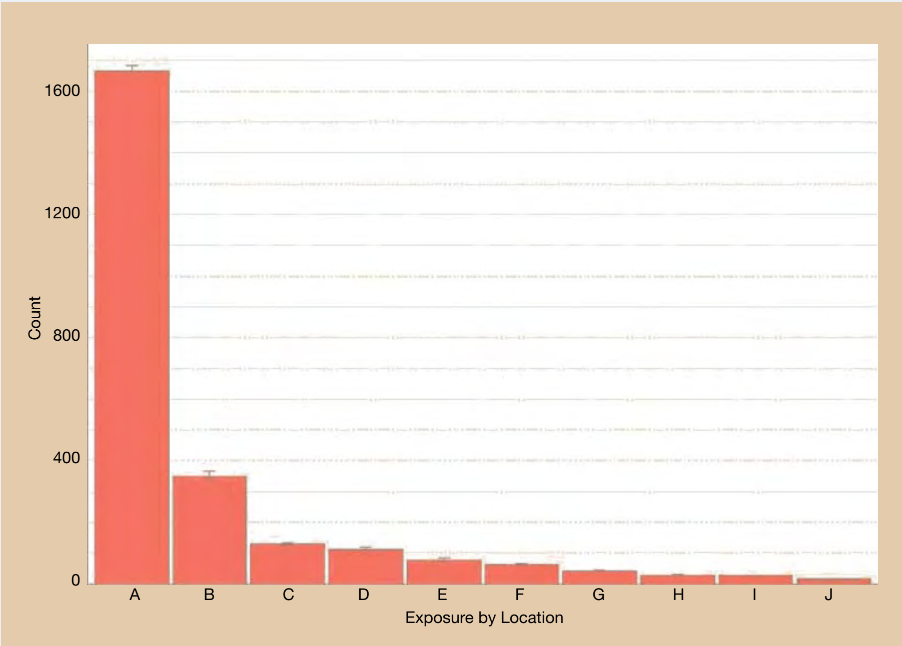

Figure 7. Number of individuals exposed to the pathogen at 10 generic locations

Disease Parameters

For simulating disease spread, CASES allows for a diverse range of pathogen parameters. These include the disease transmission rate, mortality rate, and countermeasure effectiveness. The code also enables the setting of various parameters that govern the distributions of the exposed and infectious time periods. Specifically, one can set the minimum, maximum, average, and standard deviation for these distributions using the disease parameter input file. The ability to vary these parameters enables modeling across a wide range of possible pathogen characteristics.

Additionally, because CASES returns Ro values, (the number of new infections arising from a single infected individual, or "reproductive ratio") for each simulation, disease transmission rates corresponding to Ro values from previous pandemics can be examined. This also enables determination of how Ro values vary among different ships, or even on a single ship with alternate schedules. Ultimately, the model predicts how the disease will progress through a naval ship given a particular set of disease parameters, and it lets the user determine procedures to effectively mitigate disease progression.

Results

The CASES model provides data that can help decision-makers plan responses to a pandemic. For example, Figure 7 displays the number of individuals who were exposed to the pathogen at 10 locations. This information lets decision-makers strategically allocate resources to areas where the greatest number of individuals may be exposed. End-users can also track disease spread over the course of an outbreak, which provides insight into when the greatest number of individuals may be infectious and/or potentially too sick to work. Results will change depending on the user-specified pathogen parameters and countermeasure effectiveness.

For more information and related reading materials, contact Lynn Cooper at lynncooper@mitre.org, 703-983-1731.

Using Agent-based Approaches to Improve Disaster Response Preparedness

Jennifer Mathieu, John James, Paula Mahoney, and Jeanne Fandozzi

Ensuring that government responds effectively when disaster strikes is a critical challenge facing our nation. For example, it is very difficult to test response plans at the national level. Pandemic influenza is one such disaster that if not adequately planned for, could have a devastating impact.

With employee absenteeism estimated to reach 40 percent (see page 7), far-reaching consequences will be experienced including staffing of medical facilities, power grid maintenance, communications infrastructure, as well as food and commodity distribution. Unusual patterns and high demand on communication systems may lead to overload conditions as well.

Evaluating Planned Responses

Planning for public health emergencies, such as pandemic influenza, must take place from the local to national level. Currently, there are limited means available at the national level to evaluate what is being planned at state and local levels. All levels of disaster response operational plans should be evaluated now. This includes estimating the effectiveness of planned roles, responsibilities, and procedures; determining gaps in responsibilities and risk assessment; identifying pandemic affected critical infrastructure; and performing trade-off analyses with limited resources and personnel.

To address this critical need, MITRE is conducting ongoing research to develop an end-to-end process model for evaluating planned responses for pandemic influenza at local, state, and national levels. Information, patient, and materiel flows will be represented in a discrete event simulation model to evaluate outcomes for specific scenarios. We will obtain operational procedures and scenario information from subject matter experts, government documentation, and available literature.

> **PLANNING FOR PUBLIC HEALTH EMERGENCIES MUST TAKE PLACE FROM THE LOCAL TO THE NATIONAL LEVEL.**

Our methodology uses a standard systems engineering approach. However due to the broad scope of the effort (i.e., local to national), as well as the impacted infrastructure (i.e., medical facilities, power grid maintenance, etc.), additional research is required. This includes:

- Level of detail appropriate for the model
- Performance measures for the model
- Validation of the model
- Utilizing process models as a modeling integration framework.

To date, MITRE's Epidemiological Model (see page 9), an agent-based model of disease spread, has been combined with the process model. The process model includes local to national procedures for responding to pandemic influenza, and the disease spread model is used to test various response options. How and where the pandemic starts will require different responses across the nation. Integrating these two modeling approaches is expected to improve disaster response preparedness.

Benefits

We are currently seeking collaborations with local, state, regional, and federal agencies and organizations involved in or affected by pandemic influenza response planning. Individual agencies' operational procedures for pandemic influenza will be placed into the context of the overall response.

One potential benefit of modeling operational procedures is to identify deficiencies and gaps in response plans at all levels. Tailored prototypes that can be used to evaluate preparedness at all levels will be available to collaborators in October 2008. Results of this analysis will provide valuable information to decision makers at the local to national level for planning and infrastructure analysis.

For more information on this article, contact Jennifer Mathieu at jmathieu@mitre.org or 781-271-8672.

ABM and the Validation Hurdle - An Illustration

Alfred Brandstein and Paul J. Wehner

Agent-based Models (ABM)

Our approach characterizes ABM as a way of thinking that begins with the microscopic rather than the macroscopic. An ABM perspective describes a system as a set of interacting discrete entities.

Validation

Formal definitions can be found in many places, but for the purposes of this article, we will use the following:

- Verification determines if the simulation was built right
- Validation determines if the right thing was built
- Accreditation determines if it is appropriate for a particular use.

A Challenge for ABMs

Validation requires one to gauge whether or not a simplified representation of reality (i.e., a model) is correct. For many linear and physical systems, correctness is judged by accuracy. However, for complex adaptive systems, especially those involving social interactions, a different approach to validation is required. Since data stemming from the observation of human behavior is highly variable, its value and accuracy is limited and cannot be used to accurately assess the results.

Certainly one must get the physics right, but this is sometimes trivial when compared to obtaining the confidence of the customer. Often the primary measure of merit for a model becomes: "Does the model provide non-intuitive results that can be justified?" Thus, we propose applying the following two criteria to determine if a model should be deemed valid:

- Are the outcomes produced by the model believed (by the customer) to be possible (in real life)?
- Is what is believed (by the customer) to be possible (in real life) realizable in the model?

An Illustration

The remainder of the article concentrates on a specific modeling approach developed by the authors for the U.S. Customs and Border Protection's Office of Border Patrol (CBP's OBP). Emphasis will be on methodology rather than results.

MAP AWARE NON-UNIFORM AUTOMATA OFFERS THE FOLLOWING FEATURES

- Graphical user interface that provides a high degree of transparency during model creation and simulation
- Readily supports development and use in a seminar setting, allowing for more Sponsor involvement
- Initial models could be constructed in hours
- Resulting models can run on laptops
- Extensive infrastructure exists to evaluate results from this modeling language.

The customer was interested in understanding how well various combinations of infrastructure, equipment, and personnel satisfied their mission. Algorithmic approaches and top-down representations were felt to be premature, at least until the customer was confident that the methods represented his perception of the situation. The chosen agent-based approach allowed us to identify linear regions and to explore strategies that potentially could prevent chaotic situations.

Given the recognized complexity of OBP's targeted application area (i.e., the border region), the customer played an integral role in developing a highly transparent modeling construct. It was critical for those with first-hand knowledge of the border region to be involved in the process and to speak directly with local subject matter experts. To that end, OBP co-located two border patrol agents (BPAs) with the MITRE development team for approximately six months, made others available for consult, and arranged multiple visits to the nearest border area.

The Map Aware Non-uniform Automata (MANA) ABM application was selected as the modeling construct because of MITRE's familiarity with the application, its rapid development cycle, and its easy-to-use interface that facilitates sponsor involvement. ABM was a natural choice given the high degree of human and system interactions (e.g., humans moving across a human-monitored sensor field, present in the targeted border area).

Continued on page 14

The Modeling Process

The embedded BPAs served as a clear and constant window into the targeted-application area. As they described the border environment, we distilled its key properties, captured them within the modeling construct, and verified that our representation was true to their experience. By the end of the first day, a rudimentary model was running.

Over the next few weeks, the model was further refined using the seminar context until the embedded BPAs were satisfied that all of the essential interactions were incorporated. The model was then introduced to a wider audience of BPAs who observed additional interactions, which were then identified as essential, captured, and verified. The emerging model was eventually co-presented by the MITRE developer and BPAs to CBP and OBP senior staff for review and comment.

An independent University Affiliated Research Center recommended the models for accreditation with limitations, meaning the models were deemed appropriate for the current study but would need to be re-accredited if used in a different context. The accreditation can be largely attributed to the model's methodology—one that stresses sponsor and subject matter expert involvement. This heavy involvement forces the ABM to undergo a continuous verification and validation process, making the eventual move to accreditation easier.

For more information about this article, including a complete list of references, please contact Paul J. Wehner at pwehner@mitre.org or Alfred Brandstein at brandstein@mitre.org.

DAP-E: A New Model for Predicting Human Behavior

Charles Worrell

Simulations are often used to show how a system will behave or to predict how it will perform under future conditions. With some types of systems, the best way to get an accurate representation is to simulate the system's constituent parts and then allow those parts to interact and produce the resultant behavior.

For example, to simulate the impact of a fender bender on highway traffic, one would first simulate the actions of the individual drivers as they maneuver to avoid debris, render assistance, or peek at the damaged vehicles. By analyzing the maneuvers of the individual vehicles, it becomes much easier to explain how highway traffic can slow to a crawl, even when all four lanes are clear of debris.

This technique of starting an analysis by looking at the constituent parts, or agents, is referred to as agent based modeling. We developed the Decompose, Aggregate, Propagate, and Evaluate (DAP-E) simulation technique as part of a recent MITRE Sponsored Research project to describe how agents, representing human beings, reach the conclusions that drive their behavior. Systems that include people can be hard to simulate accurately because of the variability and independence exhibited by most people. The DAP-E method provides a way to represent selected social and attitudinal factors in agent based models that may allow some simulations to produce more realistic outputs.

IN THE AGGREGATE PHASE OUTPUT FROM EACH LAYER OF PERCEPTION IS USED TO REPRESENT AN AGENT'S BELIEFS ABOUT THE ENVIRONMENT.

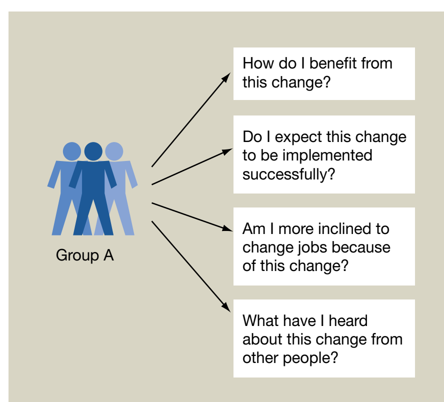

Figure 8. Decomposed Layers of Perception

The DAP-E method combines selected psychological models, a Bayesian inference technique, and a gravity model calculation to predict how changes in decisions result from changes in the agent's environment. MITRE researchers are demonstrating the DAP-E method in various domains, using agent based models to predict:

- Workforce reaction to enterprise change initiatives[1]
- Responses by undocumented immigrants to changes in U.S. immigration policy
- Public responses to environmental threats
- Reactions to changes in health-care delivery practices.

Decompose Phase

The decompose phase is an outward-looking process in which the agent forms perceptions about its environment. In this phase, the agent's awareness of the environment is broken down, or decomposed, into separate layers of perception, which can be represented by a question the agent asks itself about its circumstances. To calculate the answer, a psychological model appropriate to the domain is used.

For instance, if the enterprise change variant of this technique were used to examine the reaction of accounting staff members to the introduction of a new financial management system, the staff's perception of benefit is calculated using the Herzberg Job Motivation Model, in which measurable aspects of a work situation are used to predict whether the subject will regard the situation as positive or negative. Figure 8 shows examples of questions the accounting staff might consider while surveying their environment.

[1] The method is being examined for applications that include modeling resistance by project teams to changes in workplace procedure and the acceptance of new IT applications among stakeholder groups.

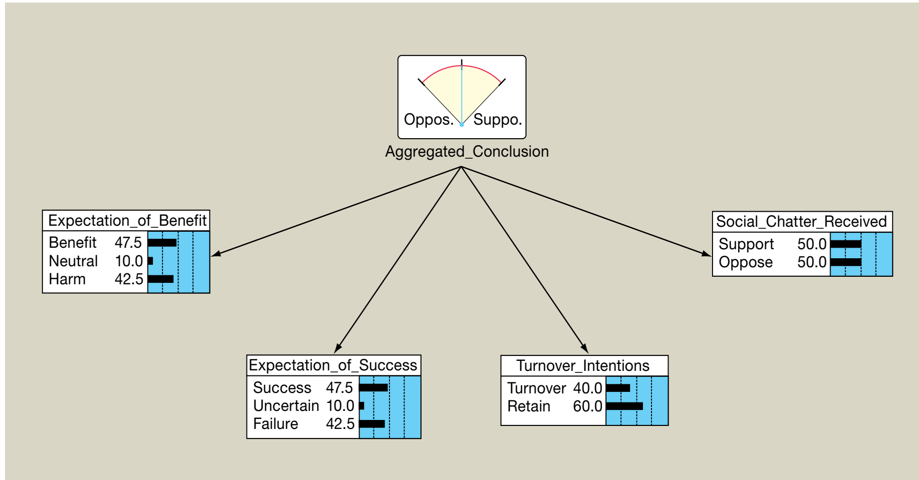

Figure 9. Aggregating Layers of Perception into a Likely Conclusion

Aggregate Phase

The aggregate phase, which is inward looking, represents the cognitive process the agent uses to draw conclusions, which are based on perceptions of the environment calculated in the decompose phase. In this phase, as shown in Figure 9, the output from each layer of perception is used to represent an agent's beliefs about the environment. These separate beliefs are combined using a Bayesian belief network to represent the agent's conclusion, which comes from a set of predetermined options and includes a confidence level. Bayesian inference techniques are used to quantify the probabilistic inferences that can be made based on the co-existence of these variables. These calculations are used to represent, for example, the accounting staff's thought process as they decide whether to support a new financial management system after surveying their environment in the decompose phase.

Propagate Phase

The propagate phase represents the influence on the agent's decisions from people who have regular contact with the agent. During this phase, depicted in Figure 10, independent conclusions from different agents are allowed to influence other agents' perceptions. As time passes in the simulation, each agent iterates through the DAP-E process, and the conclusions reached by others are used as inputs to each agent's decompose calculations. The influence of the separate opinions is determined using a set of calculations that behave more or less like a gravity

Continued on page 16

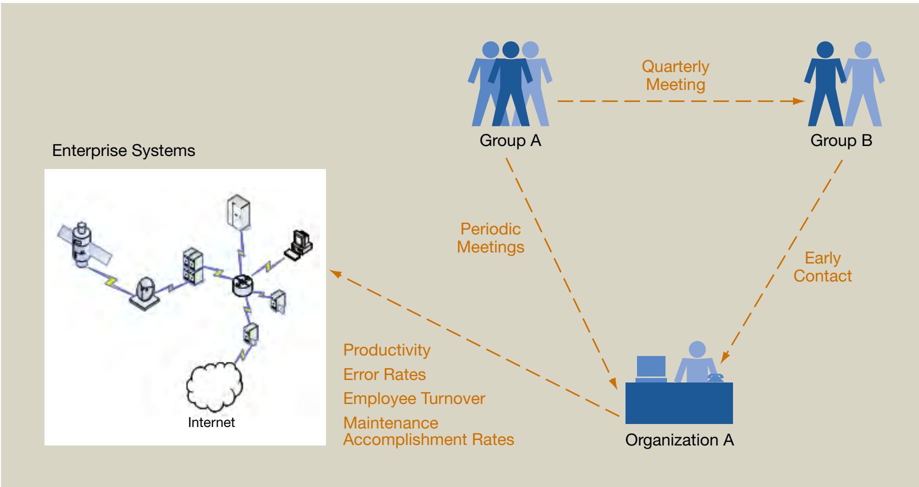

Figure 10. Propagating Opinions among Agents

model, with some agents having more influence than others. These represent the way that the accounting staff's opinions are influenced by their contacts with supervisors and colleagues from other work groups.

Evaluate Phase

In the evaluate phase, the impact of the conclusions reached in the first three phases are assessed and represented in the simulation. These can vary widely depending on the domain in which the method is applied. For example, in the enterprise change scenario, the agent's conclusion will affect its error rate or speed in performing assigned tasks. In the immigration example, the conclusion about whether to comply with U.S. immigration policy is used to reflect the migration choices of perspective immigrants.

For more information on the DAP-E Method, and related reading materials, contact Charles Worrell, cworrell@mitre.org, 703-983-1802.

About SEPO

<<

The MITRE Systems Engineering Process Office (SEPO) is a nexus for systems engineering information at MITRE. Our team brings together useful systems engineering resources, provides guidance on systems engineering processes, and participates in systems engineering activities throughout the MITRE Corporation.

SEPO works in collaboration with MITRE's Office of the Corporate Chief Engineer, the Enterprise Systems Engineering Office, and the Advanced Systems Engineering Center (ASEC).

Systems engineering resources are available through the SEPO Library, which contains a broad spectrum of information and knowledge on topics such as acquisition, systems engineering, software engineering, decision support, and process management. You can search the SEPO Library by using MITRE Google.

In addition, SEPO offers systems engineering guidance through SEPO Toolkits, including Risk Management, Requirements Process, Configuration Management, Program Assessment, and Enterprise-Node-Program Management. Guidance in other areas, such as Software Engineering, Acquisition, or the Capability Maturity Model Integration (CMMI) Process, is available from the SEPO Toolkits Team at cmmi@mitre.org.

For more information about SEPO or this newsletter, please contact Brian White, SEPO Director, at SEPO@mitre.org or 781-271-8218.

Collaborations is a publication of the Systems Engineering Process Office (SEPO). Additional copies are available.

The MITRE Corporation
202 Burlington Rd., MS M380
Bedford, MA 01730-1420

www.mitre.org/work/sepo

Jennifer Mathieu, Brian Tivnan and Matthew Koehler, guest editors
Wendy Swirnoff, editor