Principles for Effectively Representing Heterogeneous Populations in Multi-Agent Simulations

Daniel T. Maxwell[1], Kathleen M. Carley[2]

[1] 1595 Old Gallows Road, Suite 215 Vienna, VA
dmaxwell@innovativedecisions.com

[2] Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA
kathleen.carley@cs.cmu.edu

Abstract. Multi-agent dynamic-network ssimulations are emerging as a powerful technique for reasoning about complex socio-cultural systems at sufficient fidelity that they can support policy development. Within these models the way in which the agents are modeled and the fidelity of the system are critical. Basic principles guiding the development and use of these models to support policy development are described.

1 Introduction

Understanding and predicting human behavior, particularly group behavior, requires understanding and reasoning about complex systems. Examples of such systems are nation state stability, belief formation and change within societies, and the spread of infectious diseases. There are many reasons why socio-cultural systems and the behaviors that emerge from them are complex; e.g., heterogeneous populations, multiple networks connecting the members of these populations, and learning and adaptation at both the individual level and the network level.

Historically, the types of models have been applied to these complex systems have not been adequate to capture and so reason about the core sources of complexity. For example, large heterogeneous populations have been represented using deterministic macroscopic models or very simple agent based simulations with sparse representations of their information state and correspondingly simple decision logics for the agents. Information diffusion in these types of models tends to be represented using epidemiological kinds of models and either random or very simple social structures. While these models have served the community well in the preceding decade or so, recent improvements in both the social and the simulation sciences provide the ability to meaningfully improve the fidelity of information diffusion and decision making behavior in simulated populations.

Complex systems, particularly socio-cultural systems, can be most usefully understood through modeling. Multi-agent (or agent based) simulations are rapidly emerging as an extremely popular tool in this area. Applications range from simulations of colonies of ants, to networks of computers, to abstract and relatively high fidelity human populations. The types of human populations that are often represented in multi-agent simulations include commercial and social organizations, geographic neighborhoods, cities and regions of countries (Batty, 2005). Now even large geographically dispersed virtual groups or organizations that stay connected using modern communications and computing capabilities are important to consider and can be simulated.

The rapidly growing popularity of agent-based simulation is due, in part, to advances in the computer and computational sciences that make simulating large numbers of agents possible, the development of simulation scenarios easier, and visualization of simulation results more intuitive (Samuelson & Macal, 2006). Additionally, the "fit" these simulations have with the social and organizational science theories that are also being modeled are contributing to the growing popularity (Carley, 2001). While the technical advances make this type of simulation more accessible and salable, they do not necessarily consider adequately the underlying social, psychological, and analytical sciences that are essential to producing truly credible and defensible results. This knowledge gap presents a risk for some customers of the results of these simulations because the results, while either intuitive or able to be explained by some plausible story, are not well grounded theoretically. In the case of policy analyses these plausible stories could lead decision makers to choose ill-advised courses of action with potentially catastrophic consequences in terms of life and national or organizational treasure. As a very simple example of how this can occur consider an Army analysis supported by one of the authors that was trying to identify the number of trucks that the U.S. Army needed to buy. Some simulations were applied that represented a set of trucks moving supplies around a virtual theater of operations. The cargo on each of the trucks was limited by weight. Later it was learned that most truck cargo fills out the cargo space long before it reaches the weight limit. Consequently, using only one variable (weight) instead of two (weight and cube) the analysts underestimated the number of trucks required by a few hundred. Fortunately, the discrepancy was identified and the estimates were redone before the Army experienced a real truck shortage.

One type of agent-based models that are very usable for policy development is multi-agent dynamic-network models. Multi-agent models enable group, social and cultural behavior to emerge as a result of the morass of actions by social agents. Dynamic network models embedded in the simulations enable the pattern of interactions among social agents to influence and be influenced by the actions of these social agents. Together, in a multi-agent dynamic-network model the social agents act, interact, and learn. This is accomplished in a world where their behavior is constrained and enabled not just by their physical position, but by their social and cultural position in the set of networks. These networks connect individuals and groups and through which information, resources, and disease spread.

In contrast to earlier models, in the new multi-agent dynamic-network simulations not only are agents and their logic more realistic, these agents act within a social and geographical landscape that bears a higher correspondence to the real world. Information and beliefs diffuse through social networks embedded in demographic, geographical, and technical realities and these social networks change and evolve as the diffusion of information and beliefs results in changes in agent behavior.

Consequently, this new class of model, the multi-agent dynamic-network model, can be used to make more accurate predictions about the range of possible futures and consequently to study a wider range of policy issues and social activities.

Our goal in this paper is to lay out a set of principles that are grounded in the underlying social and modeling sciences that will help analysts and simulation developers to implement simulations of predominantly human populations that are well suited to their intended purpose. First, we present a set of observations from recent research by the authors in multi-agent and multi-agent dynamic-network simulation, followed by a set of principles that should be considered when designing agents to represent a large heterogeneous population in these simulations. How these agents are designed is of critical importance because agent design impacts a variety of factors including: the tradeoff between model fidelity and run time, the level and type of validation possible, and the type of virtual experiments needed to assess model outputs. In particular we will be emphasizing the representation of knowledge and beliefs in the agents, how the population of agents is described, how they interact, how information flows, and how they make decisions. To accomplish our goals we will first discuss some of the reasons we model populations and why simulation is important to understanding their behavior. We will then explore the sources of complexity in populations, followed by some examples of the current state-of-the practice in agent based simulation. We will then present a set of recommended considerations that analysts and simulation builders can use when designing multi-agent simulations or conducting analyses that apply these simulations. And, finally, we will offer a set of conclusions and some recommendations for continued research.

2 Why Use Computer Simulation

Human populations are "complex adaptive systems" (CAS). Some key characteristics of Complex Adaptive Systems that have been identified in the literature are (Dooley, 1997; Morel et al, 1998):

- Order is emergent as opposed to predetermined.
- System history is irreversible.
- System futures are often unpredictable.
- Large number of interacting parts
- Nonlinear behavior (Individually and collectively)

Human populations clearly possess all five of the characteristics identified above. This becomes even clearer when one considers how dynamic the environment in which we live has become. Complex adaptive systems cannot be understood by "just thinking about it." Rather, formal modeling techniques are needed, particularly, computer simulation models.

There are many reasons to simulate populations. Three reasons we will discuss in this paper are training, scientific inquiry, and policy support. Each of these reasons has some special purpose for investing the energy necessary to build a simulation. For example, a policy development simulation is built and applied to gain insight into some set of questions we have about how a population is likely to behave or change over time for the purpose of evaluating some set of possible interventions to identify the best course of action. Computer based simulation is an appropriate, ethical, and cost effective way of understanding the space of possibilities that are likely to emerge when various critical events occur or policies are enacted. For example, imagine that we wanted to understand whether school closures would be effective in inhibiting the spread of a pandemic influenza. Simulation actually let's us examine, in a virtual world, the probable response of populations and the spread of the influenza to interventions like school closures given flu's with different mortality and infectivity characteristics.

The overall reason that it is important to understand the underlying reasons for developing a simulation is that this knowledge is absolutely critical to identifying the fidelity that is required for achieving the desired goals. (Maxwell & Loerch, 2007) For example, in the case of a training application of a simulation the goal is to improve the proficiency of the training subjects. In most cases where multi-agent simulations are applied the target audience is a staff or response team and the specific goal is to help them work more effectively as a team. Consequently, a reasonably realistic set of results that provide a context for their interaction are sufficient for their purposes. In the case of social science research multi-agent models can be used generate hypotheses and extend social science theory (Carley & Newell, 1994; Davis, J., et al., 2007). In some instances, very simple models can be used effectively to explore a hypothesis. For example, Schelling (1971) used a simple grid model with green and red agents and just the concept of tolerance to explore how segregation occurs. Even without a concrete connection to time on the calendar, or explicit representation of agent interactions Schelling was able to explore how segregation occurs in cities. Social science research that emphasizes the diffusion of information and innovation as a central issue has been shown to require additional fidelity in the representation of peoples' (agents') knowledge to achieve meaningful results. Carley (1999) demonstrates that the social network (who talks to whom) is intertwined with the individuals cognitive picture (what they know and how they think) as well as an individuals transactive memory (perception about who knows what). This complexity implies that in order to meaningfully develop and explore hypotheses relating to information diffusion all of these concepts must be represented explicitly to gain insight into the friction those inconsistencies in these different pictures and ground truth might cause.

But why use multi-agent dynamic network simulations? Historically, the representation of population changes was accomplished using deterministic equation based computational models. For example Helbig (1992) applies an equation based fluid dynamics model to represent the movement of pedestrians around a city. Systems Dynamics models, such as Forrester's (1971) WorldII model representing global economic activity and population change, use differential equations to effect the changes in model variables over time. Mathematical models have been used by epidemiologists for decades to help researchers understand how diseases spread around populations (Bailey, 1975). These epidemiological models have even been extended and applied represent to other phenomenon like the spread of computer viruses (Kephart & White, 2001). All of these models have two properties that limit their usefulness for understanding information dependent complex adaptive systems. First, they assert a top down structure that is either static or changes using a set of centrally identified and controlled rules. Population changes occur within the constraints of the specified structure, completely limiting the ability of lower level organizations and individual entities to adapt and evolve. Consequently, the emergent responses of diverse agents resulting from adaptations in their behavior do not occur. And second, these historical approaches either lack or have a very limited representation of individual agents, the information flows among members of the population, and the effects the information has on agent behavior. This limitation relates especially to the effects multiple sources of conflicting information can have on the population. Because of this, it is difficult, if not impossible, to investigate the impact of diverse technologies and message content on the diffusion of information and the consequent change in knowledge and beliefs around the population. These characteristics limit the model's ability to sufficiently represent and explore more complex phenomenon that are based on the diffusion of knowledge and beliefs.

Policy analysis and decision support applications present perhaps the most demanding requirements for simulation fidelity and validity. (Harrison, J., et al., 2007) One reason for this is that the consumers of simulation results are often neither scientists nor analysts. Rather the consumers are operational staff and decision makers who want to know what the simulations' results tell them about their operational challenges and what they should do about them. So, there is a need for interpretability of the simulation results and a need to provide simple causal explanations of the results. This is not transparency in the sense that the exact workings and all nuances of the model need to be explained; rather what is needed are simplified explanations that get the core concepts right and communicate them clearly. Additionally, policy analysis and decision support are accomplished to support real dynamic situations. Concepts like geographic position, time, and agent behavior need to be connected to a real map, a real calendar, and real behaviors by the people implementing the policy decision. Moreover, accountability for results often occurs along very tightly constrained timelines as well as from multiple sets of stakeholders with different perspectives on the situation under study.

Multi-agent dynamic-network models afford yet another level of realism. The key limitation of standard multi-agent models is that the agent's behavior is constrained by their physical position in a grid and/or ability to move in this grid. The agents are acting in a pseudo physical space. In contrast, in a dynamic-network model the agents are positioned in a socio-cultural space that evolves as they learn and interact.

Multi-agent simulation systems make three key contributions to scientists, analysts, and decision makers. First, the development of the model helps the participants understand the relationships which come together to effect complex behavior. The disciplined process that simply building the model requires often lays bare relationships that may not have been evident before. (Sterman, 2000). Second, the model itself supports detailed analysis and enables more systematic evaluation of effects in a way that supports both explanation and forecasting. Because the patterns of behavior that emerge as a result of second and higher order interactions are grounded in a well specified set of first order relationships and behaviors it is often much easier to see and understand causal chains of reasoning that one might not have otherwise been visible. Epstein and Axtell's (1996) work using the very simple model Sugarscape stimulated a whole collection of more detailed experiments. And, the United States Marine Corps Project Albert simulations (Horne, 2001) provided simple insights that were the genesis for more detailed that the Marine Corps used to develop military doctrine. Third, because multi-agent simulations can be used to conduct virtual experiments, it is easier examine a broad range of interventions under diverse socio-demographic conditions. Gilbert (2008) points out that when one experiments on social systems "isolation is generally impossible, and treating one system while not treating the control is often ethically undesirable". Simulation allows researchers to virtually experiment on the social system without facing these design and ethical issues. Policy analyses also have the same types of issues. We have already identified that complex systems are history dependent. This means that there are no "do-overs". Once one intervenes into a social population; a multi-agent simulation can be used to engage in a series of virtual experiments for the purpose of exploring hypotheses or to conduct "what-if" analysis across a series of possible interventions that are being considered and thereby support planning.

3 The Sources of Complexity

Every journalist, and every author, is taught that to describe a scenario they need to describe the who, what, where, why, and how and when (Figure 1). Each of these factors is a source of complexity. But from a systems perspective, it is not just these entities, but also the networks of relations within and among each of these areas that contributes to complexity. Complex socio-cultural systems can be usefully represented as the set of specific nodes that populate the who, what, where why, the relations among them, and the changes in these relationships over time (the when). This representation scheme is known as the meta-network (Carley, 1999). At a practical level the ideas underlying meta networks are seen elsewhere. For example a similar concept, under the rubric "Generalized Network model" (Clark, 2006), appears in the intelligence analysis literature.

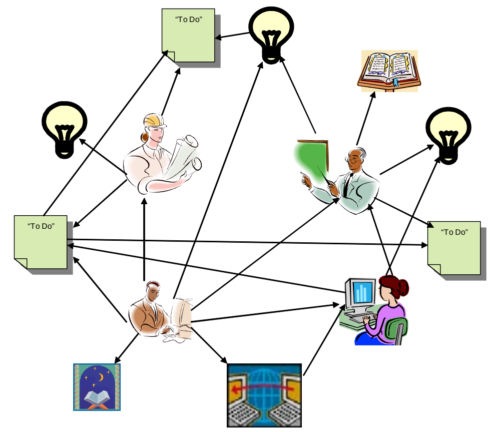

Fig. 1. The Sources of Complexity

All behavior can be described in terms of this meta-network. People interact with each other and their environment every day to work, play, and socialize. These interactions can be thought of as a combination of decisions that are based on information, beliefs, and behaviors that require and use resources. These interactions may be done to accomplish some task or take part in some event that requires certain resources or knowledge. A systematic look at people, their knowledge, beliefs, access to resources, tasks and events in which they engage, the locations at which those interactions take place, and the relationships with other agents in the environment, and the networks that form as a result of their interactions begins to reveal the true complexity of the system and consequently the modeling and analysis challenge.

In designing and building a simulation model, each of the who, how, what, where, why and when need to be addressed. The level of specificity in each of these dimensions affects the fidelity of the model. The key is to be specific enough to address the question of concern and NO MORE so as to avoid unwarranted computation, storage and analysis costs. As Einstein said, “Everything should be made as simple as possible, but not one bit simpler.”

3.1 First source of complexity: The Who!

Answering the “who” question begins to develop insight into people and their behavior. It is known that the decision-making behavior of individual humans is relatively simple. Simon (1998) points out that “Human beings viewed as behaving systems, are quite simple”. In fact humans make decisions using heuristics, moral and social norms, specified protocols, and conventions. Even though individual humans are relatively simple, human populations are complex. Simon goes on to say that “the apparent complexity of our behavior is largely a reflection of the environment.” There are many centuries of descriptive data, and libraries full of studies verbally describing this complexity. Unfortunately, this verbal information is neither complete enough, nor consistent enough to provide the detail needed for constructing a computer simulation. Given that our goal is to help practitioners build simulations, we will discuss this complexity with an eye toward identifying a set of variables that can provide a sufficient foundation for a credible simulation.

People are social animals. They are embedded in all different types of social networks: they belong to families, live in neighborhoods, have jobs, and are members of organizations. These networks affect who a person interacts with, how they interact, what they discuss and exchange, and how an interaction affects the knowledge, beliefs, and behavior of the person. It is well known that people tend to interact more frequently with similar people. In fact, McPherson, Smith-Lovin and Cook (2001) cite over one hundred studies that have observed homophily in some form or another. These factors include age, gender, class, organizational membership and role, family ties, and so forth. Consequently, a large source of complexity is that interactions among people are neither smoothly nor randomly distributed across the population. The interactions and the networks change over time, just based on who they are and how they were influenced by prior interactions. Further complicating this issue is that it is known that people will extend their social networks when they require expertise they do not possess; this is especially true if that expertise is needed to achieve some goal. In such cases, people will most often seek out the person possessing the knowledge (or resource) who is most similar to them.

In addition to people, formal and informal organizations also serve as a form of a “who” in a social system. This is because organizations have a unique identity that is more than just the sum of the interactions of its members. Organizations make decisions, have behaviors, communicate information, and have resources that are often only loosely coupled to their members. This means that there will be often be differences between the behavior, knowledge and beliefs of an organization and its members. Membership in a political party is an excellent example of how these differences emerge. (Axelrod, 1984) A political party may endorse a specific position on an issue. The party communications machinery will advertise that position widely. A member of the party may not agree with that position but choose to remain a member of the party because of agreement on other issues.

The key from a modeling perspective is that when developing a model it is critical to define who the "whos" in the model will be. Are they people, organizations, media-sources, animals, or other? Each class of who and the information-processing, cognitive and emotional capabilities they are given will impact the model results (Carley-Newell,1984).

3.2 Second source of complexity: the How!

A second source of complexity is how things get done in this network of networks. Most behaviors require resources and knowledge to be executed successfully. Resources are typically physical objects such as money, vehicles, or computers. Whereas knowledge includes specialties, expertise, ideas, key concerns, and skills. Sometimes behaviors or tasks are more complicated and may require multiple types of resources and knowledge from multiple sources to accomplish. These (often complicated) dependencies influence the likelihood which tasks will be accomplished, what communications will occur, and ultimately how the system will behave and evolve over time. (Moon and Carley, 2006)

As an example, consider the knowledge and resources required to conduct insurgent bombing activities. There is a list of possible parts. There are skills necessary to assemble, place, and detonate the device. Countering these devices also requires resources for detecting activities, interdicting supply chains, and responding to immediate threats. Depending on ones analysis goal, it could be necessary to represent explicitly most of these resources on both "sides" of the model.

One special type of resource that is especially relevant in multi-agent simulations that are concerned with information flows and decision making is information and communication technology ; i.e., the set of technical resources that allows people to obtain and share information without face to face contact. This is critical because people have differential access to information and communication technology based on a host of demographic characteristics like age, location, and socio economic status. We also know that this technology is changing all of the time. These technology based networks evolve very quickly, and in fact are not "engineered". Rather their structure evolves and their effect on the population emerges based on what segments of the population have access to and adopt the technology and for what purpose. The use of cellular telephones and text messaging in particular are wonderful examples of how this pattern plays out. Over the past few years the use of text messaging has become prevalent in younger age groups; much more so than older segments of the population. This technology effect is extremely important as we do research and policy analysis.

The key from a modeling perspective is that when developing a model it is critical to define the how. In a sense this is populating the agent with knowledge and providing access to resources. At another level, this means creating processes in the model for controlling the creation, maintenance, and depreciation of knowledge and resources, tradeoffs among resources and rules for transferring these among the agents. This includes the creation of interaction and communication protocols.

3.3 Third source of complexity: the what!

The next source of complexity is the “what”; i.e., the activities, tasks, and events in which people engage. People and other actors in the environment engage in all sorts of behaviors. In some cases the behavior involves the receipt or communication of information. In others it is the execution of a specific task to achieve some specified goal. Some tasks are simple, like making a telephone call, or paying one’s tax bill. Others are significantly more complicated and may require the execution of multiple tasks that have complicated dependencies among them. For example, successfully conducting a bombing involves procuring the required materials, assembling the device, transporting it to the target’s location, finally detonating the device. A failure on any of the tasks will result in a failure to achieve the overall task goal (Moon and Carley, 2006).

In agent based models these “whats” appear in many forms. In some models these appear as the set of actions that agents can take such as passing information, filing tax returns, going to the doctor’s office. Other agent based models take a more event based approach and define whats as a series of external interventions or events such as a speech by a local opinion leader or the closure of schools.

3.4 Fourth source of complexity: the Where!

“Where” things occur also introduces complexity into a multi-agent simulation. People, resources, events are not randomly spread around the environment. Neither are they uniformly distributed around the environment. Spatial orientation is both influenced by and influences social networks, resource networks, and task networks. People with cars and money can and do travel much further than groups without those resources. Certain kinds of activities occur in urban areas and others are more likely in rural areas. (Diseases spread faster in crowded environments.) Ponds, lakes, rivers and other geo-spatial features impact where homes and businesses are built. Location also impacts the formation and dissolution of relations. For example, two people in the same location are more likely to interact or start interacting; whereas, the tendency to interact may atrophy as they move far apart or at least the communication technology used for interaction may change.

In agent based models these “where” can appear in many forms: In many models, the where is a location on a grid and movement is dictated by this grid. In other models, locations may be defined in terms of a set of places such as home, office, or school that may or may not have specific latitude-longitude coordinates. Agent actions and use of resources may depend on location. For example, a student-agent may not be able to get information from a teacher-agent unless both are at a school.

Agents may move to locations based on their beliefs, knowledge, resources or tasks. For example, agents who are ill may go to a hospital.

3.5 Fifth source of complexity: the Why!

Another factor that influences the complexity of the social landscape is “why” things occur. The why can be thought of as beliefs, attitudes, goals, or motives. What beliefs agents hold may depend on the beliefs held by others they interact with, or their knowledge or task experience. The agent’s goals may impact where the go, whom they interact with, what information they share, what tasks they engage in, and so on.

In agent based models, the why can be implemented in many ways. Some models take a very goal oriented optimization approach. The problem here is that in many situations people can not articulate their goals (e.g., as in the case of traumatic life events), or the stated goals are not the true underlying goals (e.g., as when modeling tax compliance), or there is no way of knowing the goals of the agents modeled (e.g., as in the case of terrorists). In other models, beliefs are represented using social influence models and knowledge assessment. In this case, beliefs, like knowledge and resources can flow through the social networks and transform those networks.

3.6 Sixth source of complexity: the When!

The final source of complexity is time. The order in which things occur, lags in information flows, dictated times at which key events must occur, the length of time actions take, and so on influence the behavior of the model over time.

Complex behavior in populations is very real and has many sources. We can see that one very useful way to organize our thinking about that complexity is to look at the world as is we were a journalist. Answering the key questions of who, what, when, where, why, and how go a long way toward helping us shape our thinking about complex systems. And, we will see a little later also can help us think through the design of a simulation.

4 Illustrative Agent Based Simulation Applications

There is a growing body of multi-agent simulation work. These examples highlight what the state of the art can support, what the limits of different simulations are, and provide examples of issues practitioners should consider when undertaking a simulation development project. We will describe four different simulations, each with different characteristics. Specifically, we will look at two different simulations that address the spread of an epidemic, we will look at one simulation that supports the analysis of political and military scenarios, and the fourth describes the behavior of US taxpayers in response to different interventions by the US Government’s

Internal Revenue Service (IRS). All of these simulations are flexible enough to be used for representing scenarios other than the ones discussed in the paper. The specific scenarios help flesh out the discussion of why one is simulating as well as the specifics of implementation.

4.1 BioWar (spread of infectious disease)

A multi-university team of researchers applied a Carnegie Mellon University developed simulation called BioWar (Carley, et al. 2006) to examine the impact of life threatening events on populations at the city level. For this model, 62 diseases, including all biological warfare agents and a chemical attack for diverse cities have been modeled. This model has been used to examine the spread of anthrax, smallpox, and influenza. For example, it was used to understand the relative effectiveness of different intervention strategies, given that an influenza epidemic occurred in Norfolk (Lee, et al 2008). The simulation scenario was Norfolk Virginia and contained approximately 1.6 million agents, representing explicitly the population of that metropolitan region. The virtual environment represented homes, schools, places for entertainment, and health care facilities. Agents located themselves at one of these facilities for one or more time steps, making a decision to either change location or stay with every time step. Facilities have operational schedules and the agents have behavioral rules that preclude agents from going to closed facilities. Each time step represents four hours of real time and the scenario ran for a one year period. This particular scenario and level of fidelity required a Cray XD1 super computer to execute. Each replication of the simulation took fourteen hours to complete.

The “who” in this simulation are the people that make up the population of the Norfolk metropolitan region. Each agent was given a set of socio demographic characteristics such that the demography the virtual population was consistent with the 2000 census data. Then based on socio-demographics including occupation social networks were built among the agents to reflect human networks. Additionally, 200 agents were “infected” and spread randomly around the city. This was a proxy for a planeload of people infected with influenza arriving at the airport and disbursing.

The “what” consists of a set of actions that the agents can take. Agents relocate (or stay) every four hours and choose to interact (or not) with others while they are at that location. Also, depending on the location they may engage in other activities such as buy over-the-counter medication, get diagnosed by a doctor, become infected, spread the disease, die or become well. The choice of interaction partner is probabilistically determined, based on an agent’s demographic characteristics; the more characteristics two agents share, the greater the likelihood they will interact if they are in a position to do so. When the agents interact if one of them is infected and contagious there is a chance that the influenza will be passed from one to the other. The likelihood of this occurring is based on the transmission rate of the virus. At the same time that influenza is spreading, other diseases may as well. And, though not used in this study, the simulation also allows for airborne, waterborne, and food borne infection.

Only some of the obvious “how’s” are represented explicitly. Agent’s do have a set of knowledge, including knowledge about their visible symptoms and they use that knowledge to make decisions. Agents do get drugs, over-the-counter and prescription, and those may impact symptoms and diseases. Other general resources such as transportation vehicles are not modeled as the four hour time tick was of greater granularity than travel time. There is even a facility in the program whereby the agents know that a medical alert has been called, or the schools are closed and the agents change their behavior based on that knowledge.

The "where" in the simulated city consisted of a number of specific locations found in Norfolk including homes, workplaces, schools, pharmacies, doctor's offices, emergency rooms, stadiums, theaters, stores, restaurants, universities, and military facilities. Locations could be open or closed, depending on the day, time of day, and weather. Additionally, inclement weather or interventions could lead to school closures. The geographic, location, and weather data came from U.S. Census Bureau reports on cartographic boundaries for schools, Metropolitan Statistical Area (MSA) boundaries and business patterns, and National Oceanic and Atmospheric Administration (NOAA) Climate Data. (Carley, et al., 2004)

As indicated earlier, "why" people interact is influenced most heavily by homophily, their visible medical symptoms, and their location, and the "when" is every four hours, giving agents six possible opportunities to interact in a 24 hour period.

4.2 EpiSims (Smallpox diffusion)

Episims was originally developed at Los Alamos National Laboratory, and is now from Virginia Polytechnica Institute, for the purpose of exploring the diffusion of disease around a city or other geographical region. In the application discussed here the specific interest is the spread of smallpox around a virtual Portland Oregon. Similar to the Norfolk example researchers are interested in evaluating the differential effects of different intervention strategies and scenario assumptions, given an outbreak of Smallpox in the region. (Eubank, et al., 2004) The virtual Portland was represented using approximately 1.5 million agents and 180,000 specific locations and used the minute by minute movement around an underlying transportation network for a simulated period of 100 days to stimulate the contact that caused diffusion of the disease. The underlying simulation engine called TransSim runs on a bank of 128 networked personal computer class machines. (Los Alamos, 2008).

The "who" in this scenario consists of approximately 1.5 million agents living, working, or transiting the city of Portland Oregon. Each agent possesses a set of socio-demographic characteristics derived from the census that are used to provide the agents with differential location and mobility. The networks of connections among these people are very stylized and canonical "small world" networks.

The "how" people interact is based on the underlying transportation simulation. An agent's underlying socio-demographic characteristics are assumed to provide them with differential mobility on the network. Additionally, some agents have randomly assigned characteristics that give them even greater mobility on the network to simulate those people whose employment or lifestyle cause them to move around more broadly.

The "where" are 180,000 specific locations. The locations do not have their purpose represented explicitly. But they do have a maximum capacity which is distributed in a “scale free” fashion across all possible locations, causing agents to gravitate toward high capacity locations, like shopping malls. The belief is that this distribution is similar to how people actually behave, moving from home to public locations and back.

The “what” in this simulation is largely movement, with possible contact and infection based on co-location for periods of time. If people are at the same location for more than one hour of simulated time with an infected and contagious person, then there is some probability that the agent will contract the virus. In the scenarios aerosolized smallpox was introduced at indoors busy locations over several hours, infecting approximately 1,000 people to seed the epidemic.

The “why” for interaction is implicit in this gravitation to dense locations. The when is very highly resolved (second by second) movement data, with possible contacts occurring when two agents are in the same location for more than one hour.

4.3 Construct (Taxpayer behavior)

This simulation effort applies a Carnegie Mellon University developed multi-agent simulation called CONSTRUCT (Carley, 1991; Schrieber & Carley, 2004) to the challenge of helping the US Internal Revenue Service (IRS) identify cost effective portfolios of services, advertisements, and interventions that will encourage the US population to voluntarily meet their tax obligations. To accomplish this, the research team developed agent populations of a few thousand agents that are representative samples of the population demographics in multiple US cities. Additionally, the team developed explicit representations of tax filing knowledge, tax payer beliefs, and IRS communications programs that are designed to educate and assist the tax paying population. Simulated scenarios were run for a virtual year (sometimes two) with each time step on the simulation representing one week of calendar time. The simulation was run on a 64 bit multi-processor machine, requiring approximately three hours to complete.

The “who” in this simulation consists of approximately 3,000 or more agents representing taxpayers in a US city? These agents are imbued with a set of socio-demographic characteristics, including income and their income tax filing status. Additionally, there are “Smart Agents” (Carley& Newell, 1994) that provide taxpayers information (or misinformation) serving as proxies for newspapers, radio and TV, IRS Tax Assistance Centers, and internet access points. All of the agents have some level of knowledge about taxpaying behavior, as well as, general knowledge and some transactive memory about where to go for additional knowledge on tax related topics. Additionally, taxpaying agents have beliefs about how they feel about their obligation to pay taxes and a level of risk tolerance that is independent of their beliefs.

The “how” consists of an ability (or inability) of a taxpaying agent to engage in some taxpaying related behavior both legal and illegal. This ability is determined by the degree that the tax payer matches the target audience for the behavior. For example, the US Earned Income Tax Credit (EITC) targets lower income taxpayers to provide them with a measure of tax relief. If an individual is in the qualifying income range, they have the resources in the simulation to take the credit. Other tax related behaviors may involve other factors, like number of children or geographic location. These are also thought of resources in the simulation.

The "what" consists of two behaviors. The first is a decision to interact with another agent at every time step. The decision to exchange information and beliefs with another agent is eighty percent homophily based and twenty percent based on differential expertise. When two agents interact, they exchange a subset of their knowledge and beliefs with the other agent. In some cases the knowledge is accurate, in other cases it could be misinformation. The second behavior is an annual decision to file (or not) income taxes and if the agent chooses to file an accompanying set of decisions concerning what deductions and credits to take. This decision is based on a vector of factors, including the knowledge that the agent has acquired through interaction with other agents.

The "where" in the taxpayer simulation is only treated explicitly when it is relevant for a tax related factor. For example one study, (National Taxpayer Advocate, 2007) represented explicitly the residents of the city of Hartford in an effort to replicate an experience the IRS had with taxpayers and the local government in 2004.

The "why" agents do things in this simulation is based on an explicit representation of their beliefs about tax compliance. This set of beliefs has two components. First, the agent can believe that what they're doing is either right or wrong, or some shade in between. This is important because in many cases taxpayers engage in noncompliant behavior but have been advised either unintentionally or intentionally that the behavior is legal. The other belief is that it is either OK, or not, to cheat on one's taxes. An agent will engage in a non compliant activity if and only if one of these two beliefs is consistent with noncompliance.

4.4 SEAS – (Military Operations)

Another application of agent based simulation is in support of the US Military's effort to migrate to a new operational concept called "Effects Based Operations (EBO)". In this concept the military hopes to represent and consider explicitly the broader possible set of implications of a military operation. To do this they are describing operational inputs along four dimensions Diplomatic, Information, Military, and Economic (DIME) and outcomes along six key dimensions; Political, Military, Economic, Social, Information, and Infrastructure (PMESII).

Purdue's Synthetic Environment for Analysis Simulation (SEAS) which was originally designed as an agent simulation in support of market forecasting was adapted in three ways to analyze EBO. First, it was adapted to allow for human in the loop interaction during a simulation run, second it was federated with a simulation of military operations called the Joint Warfare System (JWARS), and finally functionality was added that was more focused on DIME and PMESII concepts. (Chaturvedi, et al., 2004)

The "who" in the simulation consists of over 100,000 simulated agents that operate in the SEAS simulation environment. The agents are given properties such that the differences among agents are consistent with the areas demographics and culture. The analysis environment also includes exogenous inputs from human players representing key government leadership roles, key neutral parties (like unbiased press), and enemy organizations. There are also data exchanged with the JWARS simulation that informs SEAS about the status of military entities that are perceived to be relevant to the PMESII variables.

The "what" in the simulation is an abstract representation of ports being opened and closed, diplomatic activities at varying levels, movement of the military and population around the environment. Agents choose to interact with each other and move about in the simulation consistent with a rule set that is based on over 15 attributes consisting of features like culture, religion, and education. Human decision makers in the key roles make higher level policy decisions and then the agents interact with each other in response the changing environment.

The "when", is relatively close to a near real-time environment, allowing for visualization of agent movement in the SEAS environment, with user selectable ability to run forward in time. This keeps the simulation timing consistent with the other simulation and allows for human interaction with the simulation. A consequence of this is normally scenarios represent days, to months of simulated time from end to end.

The "where" is normally in a geographic region of variable size and resolution. In the case of the scenario described the literature it was the city Jakarta. (Chaturvedi, et al., 2004) The environmental representation includes roads, structures, ports, and traffic loads on the infrastructure.

The "how" is based on differential resources provided to the different classes of agents, and human players. Types of resources include budgets, information, humanitarian assistance and supplies.

The "why" is based on a combination of attributes of the simulated agents, including religion, Culture, and an attribute called motivation. All of these attributes are considered as part of a rule based decision engine for determining agent behavior.

5 Multi-Agent Simulation Principles

The technical foundations of Multi-agent simulation and more generally multi-agent systems are largely based on the achievements of the computational organization theory, artificial intelligence and object oriented programming communities. Volumes have been written that describe good overall technical design and programming practices for multi-agent systems (Woolridge, 2002; Gilbert & Trisch, 2005), and simulation more generally (Law and Kelton, 2000). These practices are extremely important for good overall system design, implementation, and analysis but addressing them sufficiently is beyond our scope. Rather, we are focusing on the principle considerations that are relatively unique to development and applications of simulations to complex socio-technical systems. We highlight important considerations for six parts of the development and application process for such systems. These six parts are organized in roughly a sequential order, but practitioners should be prepared to iterate as activity in one area indicates a need for adjustments in another.

- Understand the tradeoffs
- Clearly define simulation purpose
    - If you change the purpose revisit the assumptions
    - Decide whether the model will be validated
- Use good modeling practices
    - Refine the research or analysis question
    - If the model is to be validated then identify the mapping between measurable data and simulated variables
    - Clearly specify desired output measures
    - Think explicitly about uncertainty
    - Clearly document assumptions
    - Clearly document modeling risks
- Clearly specify the variables
    - Agents
    - Environment
- Clearly specify agent behaviors
    - Use Network thinking
    - Understand how change occurs
- Conduct sufficient verification and validation testing
- Conduct well structured virtual experiments
    - Good design
    - Rigorous analysis
- Clearly present results
    - Consider the audience

5.1 Use Good Modeling Practices

All simulation (and software) development projects have a common set of decisions that must be made at the outset of the project. What language to use? How to structure the development team? What development process to follow? These are extremely important initial considerations that are discussed for the interested reader in more detail in Hoover & Perry (1990) and Maxwell & Loerch (2007).

One general consideration that warrants explicit treatment for our purposes is the need to refine the research or analysis question at the beginning of the effort. We have previously discussed the need to understand the underlying reasons for developing the simulation. That addresses the first part of this challenge.

The second general consideration is to identify the key sources of uncertainty and the planned procedures for dealing with that uncertainty. Multi-agent systems are often unpredictable due to the complexity of the system. This implies a critical need for dealing with uncertainty throughout the model development and application. Uncertainty in modeling exists at two levels. The first is the uncertainty associated with a model variable being in a specific state. These uncertainties are often addressed using probability distributions over the state space. The second kind of uncertainty, often called deep uncertainty, addresses uncertainty about the structure of the model itself. (Laskey & Lehner, 1993) One way to address effectively these uncertainties is to as exhaustively as possible describe low level behaviors, and use this knowledge as means to guide further development, experimental designs, and analysis of outcomes from the simulation. Another way to address these uncertainties is to scope the model's use to the areas where there is greater certainty. A third way is to do a series of sensitivity analyses of the critical uncertainties.

5.2 Key Simulation Design Tradeoffs

Once the decision to build or use a simulation is made, there is a series of decisions that need to be made that provide context for the rest of the model development effort. In general there are two top level design decisions. First is the level of realism, or fidelity, the agents and the environment should possess. And second, is the number of agents that will be in the simulated population. Depending on the reason for developing a model and the data that is available to support the development, the practitioner can vary these two dimensions.

In general, the higher the fidelity of the simulation system the wider the range of policy issues and social activities that can be addressed by the simulation and the more detailed the policy recommendations. However, the higher the fidelity of the simulation system the longer it takes to develop, set up, and run. The higher the fidelity the greater data requirements, and the more types and quantities of data that will be generated. Thus the tradeoff is, improved fidelity (realism) can lead to improved support for the policy analyst and decision makers. There are costs associated with this increased fidelity. The simulation will cost more and take longer to develop. The simulation will require more powerful and more expensive computational resources.

In general the higher the fidelity of the simulation system the more data that can be used to validate it and the more reasonable it is to engage in validation. That said, the higher the fidelity of the simulation, the more resources that are required to do validate the model. Time, people, and money are needed to run representative scenarios, collect data and to engage in validation relevant analyses. As these costs increase, the less likely it is that the entire system will ever be validated.

In general, the higher the fidelity of the simulation system the more types of virtual experiments that can be done to explore diverse issues and the more sensitivity analyses need to be done. What this means is the higher the fidelity the more analytically relevant data of different types that is generated. This output requires more computer storage space, advanced designs of experiments, and more statistical analysis to harvest the meaningful analytic insights. In fact, the emerging high fidelity simulation systems generally generate so much output data that standard statistical packages cannot be used to analyze the results. To meet this challenge specialized search tools and data farming environments are being developed and applied. (Horne, 2001)

In general, similar to higher fidelity representations, the greater the number of agents in the virtual population, the slower the simulation will run, and the more computing power that will be required. Matching the number of agents in the virtual population with the real population significantly increases the "realism" of the simulated scenario and simplifies many issues in the design of the virtual experiments. There are, however, sampling and experimental design techniques that allow one to conduct analyses with a smaller population of agents.

5.3 Why is the Model Being Built?

In trying to decide on the appropriate level of fidelity the practitioner needs to consider not just these tradeoffs, but also the purpose for which the model is being built. Because these are descriptive models that can not and should not represent the entirety of reality one must ensure that the explicitly represented behaviors of the agents are consistent with the key parts of the system of interest and that simulated results inform the relevant research, analysis, or policy questions. The second part is to take a small step back and do some disciplined thinking about the essential reasons one cares about developing the simulation or doing the analysis in the first place. Too often, we have seen this step in the process either overlooked because everyone on a development team "knows" the goals or timelines are too short to allow for this luxury; only to find out later this shortcut was a very costly mistake. Keeney (1992) and Edwards, et al. (2007) provide some very useful thoughts on the details of how this is accomplished effectively. Developing a clear understanding of the overall goals of the process naturally leads to an improved understanding of one's specific objectives. These objectives need to be quantified and then identified as desired outputs of the simulation. That is, from the beginning, simulation design (or simulation selection if there are existing choices) should consider explicitly output measures. Again accomplishing this exercise early brings clarity to the development process that increases the likelihood of project success.

As we stated earlier, in some cases the purpose of the model may be to explain and predict social behavior with the intention of advancing sociological theory (Carley, 2001). In other cases the simulation may be informing policy level decisions in a national government or other large organization. For example, the United States Government's Internal Revenue Service (IRS) is conducting research using multi-agent simulation in an attempt to better understand the behavior of US taxpayers so they can identify a mix of messages, services, and interventions that will help US

Taxpayers to comply with their legitimate taxpaying obligations (Carley & Maxwell, 2006; National Taxpayer Advocate, 2007). Other examples of policy decision support applications are the use of multi-agent simulation to inform planning processes in support of military operations, counterterrorism operations or strategic communication efforts around the world (See Chaturverdi, et.al. (2000) and Carley, et. al. (2003) for detailed examples). Additionally, multi-agent simulations are being used in support of training activities, particularly to generate crowd behavior emergency personnel can then interact with to practice their response to different kinds of crises. (Chiva, E & Delorme, J. 2004).

5.4 Clearly Specify the Variables

There are two overall types of variables that require specification in our virtual world: those relating to the agents and those relating to the environment in which they will function. At the very top level there are two goals for these variables. First, they must be clearly defined. That is there should be no ambiguity about what the variable represents (or doesn't). And second, they must exhaustively describe the space of entities and behaviors that are critical to the topic being studied. These concepts must be defined at a level of representational resolution that is semantically consistent both internally, and with the relevant research and study questions. (Davis & Tolk, 2007) In general, if the model is to be validated, each variable in the simulation should have a real world analog and be "measurable" in the real world. Use of variables that are impossible to measure in the real world is generally a sign that the model is to be used only for illustrative purposes.

5.4.1 Specifying Agents. The fundamental building block of an agent based simulation is an agent. Agents are most often thought of as representing people. But as we saw in the example applications, such as Construct, agents can also represent organizations, companies, nation states, computer programs, news articles and other intelligent or information processing actors. Formally, an agent is a discrete entity in the simulation that has the following characteristics:

- Autonomy – It possesses the ability to function without someone or something else having direct and complete control over its behavior. These behaviors include the ability to interact with other agents, make decisions, and to accomplish some task. This does not mean that the agent can initiate interaction merely that it can engage in some information processing operation on its own.
- Knowledge – It has information. This information may be about itself, the environment, and/or other agents. In some cases this information is both accurate and sufficient to enough to enable constructive behavior. In other cases, the knowledge is incomplete, inaccurately perceived, or incorrectly processed. This knowledge may include historical, current, or mythical information. For knowledge, the agent either knows it or does not, and for some class of knowledge may have some level of expertise.
- Beliefs – In some cases agents, especially human and organizational ones, posses a special type of knowledge called beliefs. These beliefs are usually stable and focus

on fundamental concepts like right and wrong, religious conceptions, etc. Agents can hold opposing beliefs and these beliefs can be held with some degree of intensity.

- Resources – An agent can possess or have access to resources that empower them to execute some behavior or set of behaviors. Examples of resources might be money or raw materials. Agents can hold multiple instances of resources or levels of the same resource.
- Information processing capabilities – Agents may have some ability to initiate interaction, locate information, acquire, perceive, give, process, forget information. This includes sensory ability, social skills, and adaptivity based on learning.
- Sensory ability – An agent possesses the ability to collect and perceive facts about its environment. In some cases this may be the result of direct observation, in others it could be the result of communication that is received from another agent. Social skill – Agents have the ability to communicate with other agents. In this communication they can either initiate or reciprocate and share and receive knowledge the other agent(s) in the interaction. Adaptivity – Some agents have the ability to learn and modify their behavior based on what they learn or think they learn.
- Physical capability – Agents may have some ability to move themselves or objects and so put themselves in positions to acquire or provide resources, information or beliefs, or to engage in particular actions.
- Decision criteria – Agent may have some ability to make decisions. This may be the result of clearly articulated goals and plans or simple stochastic reactions as constrained by their environment, knowledge, beliefs, and capabilities.

Not all agents possess the same level of capability along any or all of these attributes. In fact these differences among these properties are a key part of what causes the population in the simulation to be heterogeneous. Systematically mapping these properties to be aligned with the task at hand is an important early task in model development. For example, when describing an agent for use in a simulation of taxpaying behavior it is unlikely that religious affiliation and beliefs are directly relevant to the analysis, so they can be set aside. On the other hand, if the simulation is trying to explore the behavior of a population as part of a counterinsurgency analysis effort, religious belief is likely a very relevant belief. The simulation design goal is to identify a minimally sufficient set of different dimensions for describing the population.

5.4.2 Representing the Environment Finding a suitable level of fidelity for the representation of the environment is another critical part of the simulation design. In modeling the principle of Occam's Razor applies. (Jefferys, W. & Berger, J., 1992). We want to use the simplest abstraction with the fewest number of variables possible. That said, the number of variables need to be sufficient to capture the characteristics of the environment that might have major influences on the simulation's outcome. A wonderful example of this is found in the extensions of Helbig's work. The initial work simulation environment consisted of a simple grid around which agents could move. Later work demonstrated that even the introduction of simple obstacles (perhaps proxies for geographical constraints or affordability constraints) significantly changed the reults (Miyao, 1978). These differences in results caused by such a small change in the environment have stimulated significant discussion about the generalizability of Helbig's results.

5.4.3 Specifying Behavior The next task in building a multi-agent simulation is to specify the set of actions or behaviors that the agent can engage in, conditions for the behavior and response to that behavior both by that agent, the environment, and other agents. We have previously discussed the sources of complexity that make this a challenging task using a journalist's perspective. We can use these dimensions to help us organize our thinking about behavior. Figure 2 organizes those same concepts into a multidimensional network that specifies the set of possible factors that might influence agent behavior and subsequently simulation results. Combining an inspection of this network with what we have learned previously about multi-agent simulations, we can see that the "who", "how", and "why" make up the set of agents and the attributes those agent possess that influence their behavior in the environment. The "what" are the things the agents do in the simulation. And the "where" is the (physical and/or virtual) location of those actions. In fact one thinks about this in object oriented programming terms, the agents and locations are objects, with the how and the why as attributes of the object or constraints on what that object allows, the behaviors are methods in the software that represent the behavior of the agents. For each agent, its environment is the collection of other agents, and the set of possible how, what, where, and whys that exist at a specific when. Each agent's perception of this meta-network is that agent's transactive memory.

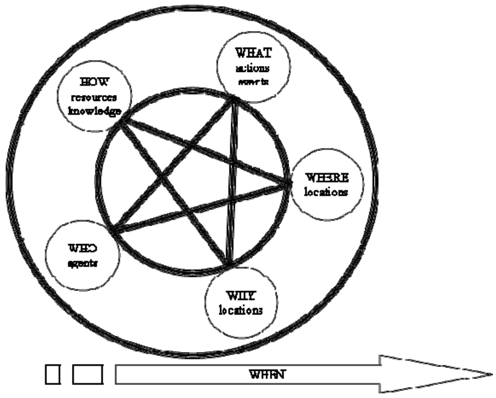

Fig. 2. The meta-network as it evolves through time

This meta-network specifies a matrix of relations, often referred to as the meta-matrix. Specifying the meta-matrix is a helpful way to specify key aspects of the design of the multi-agent simulation system. Adding another level of detail to our thinking is even more helpful for actually designing and developing a specific model. Table 2. lays out a matrix that identifies the types of networks that should be considered as part of the design. For example, the simulation requires some abstraction for "who knows who". In BioWar and Construct the agent has a set of alters that are in its sphere of influence that it normally interacts with, but based on where the agent is when or what events are ongoing the agent my engage in potentially random interactions with others outside this sphere. These networks may reflect known socio-demographic constraints as in BioWar or be specified using very stylized hypothetical structures such as the scale-free networks in EpiSims. Specifying all the relevant relationships among all of the agents in the agent population a priori would be a daunting task with many thousands of agents. A possibility is to specify a set of attributes that allow one to adjust the probability that two agents will interact based on their similarity. In the IRS example above the authors used a vector of attributes, including age, gender, race, income, marital status, and education as indicators of similarity. The more attributes the agents held in common, the higher the likelihood that they would interact and exchange information. More generally, we need to have some representation of the fact that people are embedded in different social networks, and that those networks influence who interacts.

Table 2. Network view of illustrative relationships in a multi-agent simulation at a particular point in time

	Agents	Knowledge	Resources	Actions	Events	Locations	Beliefs
Agents	Social Network	Knowledge Network	Capa-bilities Network	Activities Network	Partic-ipation Network	Physical Presence	Belief Network
Know-ledge		Information Network	Skills Network	Know-ledge Needs			Factual basis
Resources			Substi-tution Network	Resource Needs		Availability	
Actions				Workflow Network			
Events					Prece-dence Network		
Locations						Borders	
Beliefs							

The concept of differential access to information (knowledge) is further complicated when we introduce the ideas behind the knowledge network and the capabilities network. For example, there is some type of information available through internet sources. To represent this effectively in research or analysis we would likely need to create a "smart agent" that has that knowledge available for an agent to find. Then the searching agent would have to be able to access the internet (a resource) and either have knowledge that the information was available (transactive memory) or have some behavior that allowed it to search the internet. Again, one could randomly indicate that agents have access to the internet, but as we saw earlier technology use is very different depending on age, and other demographic features. This means that in cases where understanding information flows is central to the research or analysis question then a more detailed model is warranted.

In designing and developing a multi-agent simulation, the types of thought experiments and considerations described above should be conducted for all relevant networks described in Table 2. A useful technique for conducting these thought experiments is to think about what is (or might be) flowing through the network. In some cases we are looking at the transfer of material goods and services to satisfy and economic demand (Chaturvedi, A et al., 2008). In the case of material goods and services it is important to specify the set of attributes and relationships that prevent very unrealistic things from occurring. (e.g. infinite supplies of goods) In other cases we are exploring the propagation of knowledge and beliefs throughout a population. Again, in some cases randomization or cursory treatment of these concepts may be sufficient, but others will require explicit treatment. Our experience is that this requirement is especially prevalent in the representation of an agent's knowledge, beliefs, and decision making processes. This is because not all messages are created equal. As a simple example, think about a message that addresses the belief about a behavior being "right or wrong". A website, radio message, or news article will have a different impact than a parent or a cleric communicating a message. Think about situations where the messages are different, depending on the goals of the agents. One example of this might be a simulation of strategic communications in a counterinsurgency environment. Just asserting the effectiveness of some positive messages, or even positive behavior, with out simultaneously considering the effects of conflicting messages will more often than not lead a researcher or analyst to become overly optimistic about how the system will respond to stimulus.

It is true that one can not always predict what the overall changes in a complex adaptive system will look like. This does not imply that one does not need to understand the processes that produce those changes. In fact, it is absolutely essential that the basic processes be both understood, and clearly described as functions (software methods) in the simulation software. There are three facets to achieving an understanding of change processes. The first is developing an understanding about the nature of change. For example, providing a person with money will certainly increase the amount of money they have immediately available. And, the additional money makes it more likely that they will save some of it. In the simulation, the first change, receiving money, is likely best represented just by adding financial resources to the agent. In the second case, we probably want to represent the action "saving" as a decision that is a function of the additional money, but also considers other goals, knowledge, and beliefs held by the agent. This ability to locally consider multiple factors is an essential part of the power of multi-agent simulation. The second facet is to understand the rate of change. In the case of money transferring from agent to agent, the effect immediate. But it says nothing about how frequently the transfer occurs. Is it weekly or monthly, like a paycheck? Or is it annually, like a tax refund? Other types of changes occur along different timelines. For example prejudicial beliefs will likely change very slowly. Years to decades could go by before any meaningful change is seen in deeply held beliefs and engrained attitudes. Consequently, practitioners should be careful not to be overly optimistic or specific about the rates at which change happens. The final facet of change that must be understood is the mechanism that executes change. Using our financial example again, moving money from one agent to another changes their resource position. That is the nature of the change. The mechanism describes how the transfer takes place. Is it a cash transaction in person; is it a check in the mail, or an electronic fund transfer? Different mechanisms for change will have implications on what can and should be represented explicitly, or left out as peripheral to the question under study.

6 Model Verification, Validation and Testing

Depending on why the model was built, verification and validation (or V&V) activities may be an important aspect of project success with multi-agent simulations of complex socio-technical systems. If the model was built to simply illustrate a process (as in training) V&V is not warranted. In principle, the level and type of V&V depends on the maturity of the model and the uses to which it will get put (Zacharias, et al. 2008). In practice, the level and type of V&V depends heavily on available empirical data, which may not exist in sufficient quantity for statistical validation, and the resources available for validation. Verification activities answer the question "Did I build the model right?" and validation addresses the question "Did I build the right model?" There is a literature available that describes a wide range of techniques for conducting V&V (e.g. Balci, 1994; Windrum, et al. 2007).

On the verification side, good software engineering practice, version control and testing are needed. Even relatively simple multi-agent simulations, with simple local interactions can quickly become rather complicated pieces of software, with significant amounts of data. So, the entire system needs to be methodically and comprehensively checked for errors in implementation, design, and data.

In terms of validation, there are a plethora of validation challenges that face multi-agent simulations of complex socio-technical systems that simply do not arise for engineering level simulations. The range, level and types of validation techniques have exploded with the complexity of these systems. Early researchers in artificial intelligence used the idea of a Turing Test (Turing, 1950) as a test for the quality of a computational model attempting to replicate human behavior. The basic idea is that if you interact with a device and cannot tell whether the information is coming from a computer or a human then that computer model is an adequate model of the human. This is a test of the isolated individual. In multi-agent simulations we have social agents. Carley and Newell (1994) introduced the idea of a "social Turing test" that can be helpful for V&V purposes for social agents. The social Turing test is "weaker than the Turing test because it does not require confusing a computer with a person. It is stronger because it allows for plugging "in many values." The basic idea is that if you see results about a group or population generated by a device and cannot tell whether these results were gathered from a real group or population or generated by computer simulation, then the computer simulation model is an adequate model of the social milieu. This type of test would allow a researcher to explore the range of possible behaviors for an agent under many different sets of conditions to assess its validity. Moreover, if this is done systematically across the set of agents the test could also provide some insight into the reasonableness of the behavior of the population as a whole as well.

At the time Carley and Newell introduced the idea of the "social Turing test" they indicated that "carrying out such a test is well beyond the current art," as no known simulation at that time had an adequate model social agent. We note that even for the sophisticated models described here, the model social agent's are nearer but still not completely consistent with the full range of behavior expected of the model social agent.

Practitioners need to pay particular attention to the available data, not only during model development, but also during validation verification, and testing. The quality of the available supporting data supporting simulation input speaks volumes about what simulation results might be useful for. The data confirms what we know and guides what we can reasonably infer from the simulated results. Therefore, the input data should be evaluated to ensure it is relevant to the variables that are being populated. It should be reasonably available or relatively easily (and reliably) imputed from data that is available. As an example, for a recent IRS study we had census data readily available for describing the overall population of a particular community. But, the focus of the study was on a relatively small subset of the population. In order to improve the emphasis on the study population we applied a matched sampling technique and imputed an agent population that emphasized the study population, but allowed us to statistically infer the population level statistics from the matched sample. (Rubin, 2006) Additionally, some data are uncertain. These uncertainties should be clearly documented and explored using sensitivity analysis techniques to evaluate what the impact of that uncertainty might be on the simulation's results.

7 Conducting Virtual Experiments

The execution of a virtual experiment, similar to all other types of experiments, requires serious thought and some planning to effectively execute. The quality of the experiment's design and analysis can have as much influence on the usefulness of the results as the entire development effort. Moreover, the previously identified limitations on the ability to validate multi-agent simulations make it even more critical to think through carefully how the simulation is to be applied.

A first step is to select a set of independent variables that are relevant to the policy alternatives or hypotheses that are under consideration. Then for each variable a set of levels need to be chosen that are "representative" or critical for these alternatives. For example, if we are examining the impact of IRS interventions then the variables might be the presence of an intervention such as a newspaper add or letter to taxpayer. Whereas the level would be the number of days the add runs or the number of letters. These levels should cover the range of possibilities for the independent variables and so sample the entire space of feasible possibilities (e.g. no intervention at all, maximum intervention possible). The sample should also include sufficient interim possibilities to allow for the possibility of identifying trends, especially points of inflection in response surfaces over the results.

As we stated earlier dependent variables, or simulation results, in the experiment should inform the analysis of the overall effort. In the case of training, it is likely that the dependent variables will be performance data collected through observation of the human participants and their supporting tools. The simulation appropriately provides context for that interaction. In the case of scientific research and decision support the dependent variables should relate directly to the key hypothesis or hypotheses, or the fundamental objectives of the decision makers. In cases where the dependent variables are indirect proxies for the fundamental objectives the asserted relationships between the proxies and the fundamental objectives should be thought through and documented before the experiment is executed. This disciplined step will help frame the results analysis effort and increase the quality and defensibility of the conclusions.

In general, one might want to apply standard experimental design constructs such as Box-Behken. However, it should be noted that the number of variables in these models is sufficiently large and the variables themselves may not be continuous that design constructs devised for physical device simulations may not be appropriate.

Experimental designs need to consider the possibility of interaction effects across the set independent variables under study. Too often in simulation studies of all sorts independence among the independent variables is just asserted and experiment designs are then built on that assumption. It is more often the case, especially in multi-agent simulations that there are interactions among the independent variables. In some cases the interactions provide greater than additive returns (i.e. synergy) and in others the interactions can reduce the expected returns to something less than additive (destructive interaction). A better approach is to develop design for experiments that test for independence rather than assert it. Those dependencies that emerge can then be explored in some detail, either in search of an opportunity in the case of synergy or to mitigate downside risk in the case of destructive interactions.

Once the virtual experiment is run the results need to be statistically analyzed and the local response surface estimated. Often the desire to identify high coefficients of correlation and tight confidence intervals on the statistics can increase with the number of replications to infeasible levels. This is further complicated because even you if you achieve tight confidence intervals and a high R2 by increasing the number of replications, one still may not have the kind of insight that is necessary for drawing meaningful conclusions. Consequently simply judging the results using significance levels can be meaningless. To deal with this reality experimenters should do considerable sensitivity analyses (Gilbert, 2008), and look to nonparametric statistical analysis techniques and the use of data farming environments to generate sufficient results for meaningful analysis. When analyzing these results concentrate on the relative value of the beta (standardized) coefficients rather than just the significance levels.

8 Presenting Results

The use of multi-agent simulation in support of policy decision presents a special presentation challenge. Decision makers and operational staffs are very focused on reducing and where possible eliminating uncertainty and then selecting and executing the “right” course of action. Multi-agent simulation makes explicit uncertainties they have always faced, but may have necessarily assumed away to keep their planning and analysis challenges feasible. The analysts challenge in presenting multi-agent simulation results is to quickly and effectively communicate what the simulation is saying and why it is important to them.

The age old guidance “keep it simple” is extremely important in operational decision and policy analysis environments. As we indicated earlier, the results need to be transparent, not at the level of describing the complicated inner workings of the simulation; rather they should make clearly visible the cumulative if-then result for that decision of policy context. One rule of thumb is to structure the presentation to very quickly answer two questions; “So what?” and, why?” Normally, the so what is a report on visible trends in the movement of the dependent variables in response to some of the different treatments. The “why” can be much trickier in explanation because the trends, especially counterintuitive ones, emerge as a consequence of second and higher order interactions over time. A very effective approach we have used to illuminate the results is to complement the aggregate results with a couple of anecdotal descriptions of how a few agents got to the state. This can be especially helpful for counterintuitive cases.

Another important thing to consider in presenting multi-agent simulation results is in the area of uncertainty. Multi-agent simulations are stochastic, and often illuminate significant uncertainties that exist in the environment. These uncertainties should be discussed explicitly with decision makers in terms that make their significance clear and actionable whenever possible. Talk about the results as identifying the space of possibilities. Highly replicated results are those that are more probably, but not guaranteed to occur. Examples of cases in point are combinations of factors that might be seen that present a special opportunity for achieving exceptional results, or a set of circumstances that highlight a rising likelihood of bad outcomes that can be mitigated by some condition based operational action.

9 Conclusions and Recommendations for Further Research

We have seen that multi-agent simulation has been and can continue to be usefully applied for many applications. We have identified a set of things that should be considered in each of the applications as well as a set of general principles to be followed for most effective simulation development. Adherence to these principles can be a key contributor to the long run success of a multi-agent simulation development effort.

As we have indicated, multi-agent simulation is a rapidly growing and rapidly changing field. There are still a number of critical unanswered questions that will benefit from ongoing and future research. Future simulation developers and users would benefit from looking at the current state-of-the-art when they read this paper. Some areas of particular interest are:

- Validation and Verification practices for multi-agent simulations
- Multi-resolution modeling in simulations
- Statistical analysis techniques for simulation results to include data farming techniques
- Presentation techniques for communicating multi-agent simulation results

Authors

Daniel T. Maxwell is a Senior Principal at Innovative Decisions, Inc. In that capacity he is leading research efforts into Social Network Analysis and related simulation and analysis techniques for various government agencies. Dan also provides modeling and analysis support directly to senior leaders across the DoD, and

other government agencies. Dan has research interests and has published in the areas of decision analytic modeling, simulation, resource allocation, and modeling and analysis for organizational leaders. Prior to joining IDI Dan led many research and analysis efforts as a corporate officer in a private research firm and in leadership roles for a Federally Funded Research Center. Dan completed a full career in the US Army where he served as the Technical Director of the Joint Warfighting System (JWARS) Office in the Office of the Secretary of Defense (OSD). Dan has also served in the U.S. Army Concepts Analysis Agency where he provided analysis in support of many DoD studies. He holds a Ph.D. in information technology from George Mason University, an MBA from Long Island University, and a B.S. from the Rochester Institute of Technology.

Kathleen M. Carley is a Professor of Computer Science and the Director of the Center for Computational Analysis of Social and Organizational Systems (CASOS) at Carnegie Mellon University. She received her Ph.D. from Harvard in Mathematical Sociology, and her undergraduate degrees are from Massachusetts Institute of Technology. Her research combines cognitive science, social networks and computer science to address complex social and organizational problems such as belief formation in populations, counter terrorism and counter narcotics. She is the co-founding editor of Computational and Mathematical Organization Theory, the founding president of the North American Association of Computational Social and Organizational Sciences, has developed multiple multi-agent dynamic-network simulation models, and has published over 130 articles and edited multiple books in this area.

Acknowledgments

The authors would like to acknowledge the support of multiple US Government organizations in both the preparation of this paper and our research in recent years. In particular we would like to recognize the support of Ms. Patricia McGuire, Director of the Office of Policy Evaluation and Risk Analysis at the Internal Revenue Service. Her vision and persistence have made invaluable contributions to our work in simulation. Of course any errors are the responsibility of the authors alone.

References

Axelrod, R. (1984) *The Evolution of Cooperation*, Basic Books, New York.

Bailey, N. (1975), *The mathematical theory of infectious diseases and its applications*, second edition, Oxford University Press, New York.

Balci, O (1994) “Validation, Verification and testing techniques Throughout the lifecycle of a Simulation Study” *Annals of Operations Research*, vol. 23, pp 121-173

Bankes, S. (1993) **“**Exploratory Modeling for Policy Analysis**”** *Operations Research*, Vol. 41, No. 3 (May - Jun., 1993), pp. 435-449

Batty, (2005). *Cities and Complexity: Understanding Cities with Cellular Automata, Agent-Based Models, and Fractals,* MIT Press, Cambridge.

Brown, R.E. & Mazur, M.J. (2003) “IRS’s Comprehensive Approach to Compliance Measurement”, IRS White Paper, Washington D.C.

Carley, Kathleen M.1991. “A Theory of Group Stability.” *American Sociological Review,* 56(3): 331-354.

Carley, K. (2001). Computational Approaches to Sociological Theorizing. 69-84. J. Turner, ed.

Carley, K. (1999) “On the Evolution of Social and Organizational Networks” See http://www.casos.cs.cmu.edu/events/summer_institute/2001/reading_list/pdf/EvolutionofNetworks.pdf

Carley, K. et al. (2003), “Destabilizing Dynamic Covert Networks”, In *Proceedings of the 8th International Command and Control Research and Technology Symposium.* Washington D.C.

Carley KM, Fridsma DB, Casman E, et al. (2006) “BioWar: Scalable Agent-Based Model of Bioattacks”, *IEEE Transactions on Systems, Man and Cybernetics -- Part A: Systems and Humans* vol. 36, pp. 252-65.

Carley KM, Altman N, Kaminsky B, Nave D, Yahja A. (2004) BioWar: A City-Scale Multi-Agent Network Model of Weaponized Biological Attacks: CASOS Technical Report: CMU-ISRI-04-101. In: Carnegie Mellon University, Pittsburgh, PA.

Carley, K & Maxwell, D. (2006) “Understanding Taxpayer Behavior and Assessing Potential IRS Interventions Using Multi-Agent Dynamic-Network Simulation”, *Proceedings of the 2006 Internal Revenue Service Research Conference, Washington D.C. June 14-15, 2006.*

Carley, K & Newell, A. (1994) “The Nature of the Social Agent”, *Journal of Mathematical Sociology,* vol. 19, no. 4, pp. 221-262

Chaturvedi, A; Dehnke, R. & Snyder, D.. (2004) “Simulating Nonkinetic Aspects of Warfare”,” *Proceedings of the Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC), 2004.*

Chaturvedi, A.R.; Gupta, M.; Mehta, S.R.; Yue, W.T **(2000)** “Agent-based simulation approach to information warfare in the SEAS environment**”**, in System *Sciences, 2000. Proceedings of the 33rd Annual Hawaii International Conference on System Sciences.*

Chiva, E. & Delorme, S. (2004) “Performance of Motivational Command Agents in a Command Post Training Simulation”, *Conference on Behavioral Representation in Modeling and Simulation (BRIMS, 2004),* paper.

Clark, R. (2006) *Intelligence Analysis: A Target Centric Approach,* xxxxx

Davis, J. P., Eisenhardt, K. M., & Bingham, C. B. (2007). Developing theory through simulation methods. Academy of Management Review, 32, 480-499.

Davis, P. & Tolk, A. (2007) “Observations on New Developments in Composability and Multi-Resolution Modeling”, *Proceedings of the Winter Simulation Conference 2007,* Washington D.C.

Dooley, K. (1997): "A Complex Adaptive Systems Model of Organization Change," Nonlinear Dynamics, Psychology, & Life Science, Vol. 1, No. 1, p. 69-97

Edwards, W. ; Miles, R & Von Winterfeldt, D. (2007) *Advancies in Decision Analysis From Foundations to Applications,* Cambridge University Press, Cambridge.

Epstein, J. & Axtell, R. (1996) *Growing Artificial Societies: Social Science From the Bottom Up,* Brookings Institution Press, Washington D.C.

Eubank, S.; Guclu, H.; Kumar, V.; et al. (2004) “Modeling Disease Outbreaks in realistic Urban Social Networks”, *Science,* Vol. 429, 13 May 2004.

Forester, J. W. (1971) *World Dynamics,* MIT Press, Cambridge, MA.

Gilbert, N. (2008) *Agent-Based Models,* SAGE Publications, Thousand Oaks, CA.

Gilbert, N. & Troisch, K. (2005) *Simulation for the Social Scientist,* Open University Press, Berkshire.

Harrison, J. R., Lin, Z., Carroll, G. R., & Carley, K. M. (2007). "Simulation modeling in Organizational and Management Research" . *Academy of Management Review*, 32, 1229-1245.
Helbig, D. (1992) "A Fluid Dynamics Model for Movement of Pedestrians", *Complex Systems,* vol. 6 pp. 391-359
Horne, G. (2001) *Maneuver Warfare Science,* US GPO, Washington D.C.
Jefferys, W. & Berger, J. (1992) "Ockham's Razor and Bayesian Analysis", *American Scientist,* vol. 80, pp 64-72
Kaufer, D. & Carley, K. (1993) "Comunication at a Distance: The nfluence of Print on Sociocultural and Organizational Change",
Keeney, R. (1992) *Value Focused Thinking: A Path to creative Decisionmaking,* Harvard University Press, Cambridge.
Kephart, J. and White, S. (1991) ``Directed-Graph Epidemiological Models of Computer Viruses," *Proceedings of the 1991 IEEE Computer Society Symposium on Research in Security and Privacy,* Oakland, California, May 20-22, pp. 343-359.
Laskey, K. & Lehner, P. (1993) "Metareasoning and the Problem of Small Worlds", unpublished manuscript.
Law, A. M. and W. D. Kelton (2000), *Simulation Modeling and Analysis,* third edition, McGraw-Hill, NY.
Lee, B. et al. (2008) "Virtual Epidemic in a Virtual City: Simulating the Spread of Influenza in a United States Metropolitan Area" Unpublished manuscript.
Los Alamos (2008), Information found at: http://transims.tsasa.lanl.gov/
Maxwell, D. & Loech, A. (2007) "Executive's Guide to "Practical Computer Models", *Public Manager,* Vol. 36. no. 3, Fall 2007.
McPherson, M., L. Smith-Lovin, and J. Cook. (2001). Birds of a Feather: Homophily in Social Networks. *Annual Review of Sociology*. 27:415-44.
Miyao, K. (1978) "Dynamic Instability of a Mixed City in the Presence of Neighborhood Externalities" *The American Economic Review*, Vol. 68, No. 3, pp. 454-463.
Moon, I. & Carley, K. (2006) "Testing the Robustness of Team Structures with Social Simulation", *Proceedings of the WSC 2006"*, Kyoto August 21-25.
Morel, B. & Ramanajuan, R (1998) "Through the Lokking Glass of Complexity: The Dynamics of Organizations as Adaptive and Evolving Systems", Organizational Science,
National Taxpayer Advocate (2007) *National Taxpayer Advocate 2007 Report to Congress,* USGPO, Washington D.C.
Parunak, H.V.D. , Savit R., Riolo R.L. (1998) "Agent-Based Modeling vs. Equation-Based Modeling: A Case Study and User's Guide" In. *Multi-Agent Systems and Agent Based Simulation*, Springer-Verlag, New York.
Rubin, D. (2006) *Matched Sampling for Causal Effects, Cambridge University Press, Cambridge.*
Samuelson, D.A. and Macal, C.M. (2006) "Agent-Based Simulation Comes of Age", *OR/MS Today*, Vol. 33 no.4, pp. 34-38
Schelling, T. (1971) "Dynamic Models of Segregation", *Journal of Mathematical Sociology,* Vol 1. pp. 143-146.
Schreiber, Craig and Kathleen M. Carley, 2004. Construct - A Multi-agent Network Model for the Co-evolution of Agents and Socio-cultural Environments Carnegie Mellon University, School of Computer Science, Institute for Software Research International, Technical Report CMU-ISRI-04-109.
Simon, H. (1998), *The Sciences of the Artificial",* MIT Press, Cambridge.
Sterman, J. (2000) *Business Dynamics; Systes Thinking and Modeling for a Complex Wolrd.* McGraw-Hill, New York.
Turing, A. (1950) "Computing Machinery and Intelligence", *Mind, vol. 59, pp. 433-460.*
Waldrop, M. M. (1992) *Complexity: The Emerging Science at the Edge of Order and Chaos*

Windrum, P.; Fagiolo, G. & Moneta, A. (2007) "Empirical Validation of Agent-Based Models: Alternatives and Perspectives", *Journal of the Artificial Societies and Social Simulation*, vol. 10, no.2,8.

Woolridge, M. (2002), *An Introduction to Multi-Agent Systems*, John Wiley & Sons, London.

Zacharias, G.: MacMillan, J; Van Hemel, s. (ed.) (2008) *Behavioral Modeling and Simulation: from Individuals to Societies,* National Academies Press, Washington, D.C.