

A Large Scale, High Resolution Agent-Based Insurgency Model

Running Head: A Large Scale, High Resolution Agent-Based
Insurgency Model

Innovation Program Report: 026-08 E

Lucas A. Overbey^{1*}, Bradley Mitchell^{1*}, Samuel Yaryan², and Kenneth
McCullough¹

¹*Advanced Technology Transition Branch (Code 56150), Information Warfare
Exploitation Systems Engineering Division, Intelligence Surveillance and
Reconnaissance / Information Operations, Space and Naval Warfare Systems
Center Atlantic, P.O. Box 190022, N. Charleston, SC 29419-9022 USA*

²*Information Warfare Exploitation Systems Engineering Division (Code 56100),
Intelligence Surveillance and Reconnaissance / Information Operations, Space and
Naval Warfare Systems Center Atlantic, P.O. Box 190022, N. Charleston, SC
29419-9022 USA*

Corresponding Author: Lucas A. Overbey, Space and Naval Warfare Systems Center Atlantic,
P.O. Box 190022, N. Charleston, SC 29419-9022 USA, E-mail: lucas.overbey@navy.mil,
Phone: 843-218-3668

*The first two authors contributed equally to this work.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 30 SEP 2013		2. REPORT TYPE		3. DATES COVERED	
4. TITLE AND SUBTITLE A Large Scale, High Resolution Agent-Based Insurgency Model				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Lucas Overbey; Bradley Mitchell; Samuel Yaryan; Kenneth McCullough				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Space and Naval Warfare Systems Center Atlantic,P.O. Box 190022,N. Charleston,SC,29419-9022				8. PERFORMING ORGANIZATION REPORT NUMBER 026-08-E	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT Recent years have seen a large growth in research developed aimed at modeling the intricate social-cultural climate in which we operate. In particular, agent-based models have garnered considerable attention for their ability to simulate heterogeneous actors operating in complex environments. This paper introduces an agent-based insurgency model capable of simulating hundreds of thousands of agents with complex cognitions and behaviors. Unlike previous agent-based models of civil violence, this work includes the use of a hidden Markov process for simulating agent decision-making, the ability to actively or passively aid insurgency or counterinsurgency efforts, and the incorporation of environmental constraints on agent actions. To significantly improve performance and scale the simulation to hundreds of thousands of agents, the distributed processing capability of a graphics processing unit was utilized. With this model, we are able to show several emergent behaviors such as localized outbursts of insurgent activity, dynamics representative of punctuated equilibrium, and tipping points leading to major insurgent activities. In addition, data based on actual events in Ramadi, Iraq in between 2003 and 2008 was generated and correlated with simulation results to demonstrate how these models can portray realistic insurgent environments.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT 3	18. NUMBER OF PAGES 28	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Abstract

Concurrently with Department of Defense interests and requirements, recent years have seen the growth of technologies developed to model the complex political, military, economic, social, infrastructure, and informational systems in which we operate. In particular, agent-based models have garnered considerable attention for their ability to simulate large numbers of heterogeneous actors operating in complex environments. This report introduces an agent-based insurgency model capable of simulating hundreds of thousands of agents with complex cognitions and behaviors. Novel innovations for this model include the use of a hidden Markov process for simulating agent decision-making, the ability to actively or passively aid insurgency or counterinsurgency efforts, and the incorporation of environmental constraints on agent actions. To significantly improve performance, the distributed processing capability of a graphics processing unit was utilized. With this model, we are able to show several emergent behaviors such as localized outbursts of insurgent activity, dynamics representative of punctuated equilibrium, and tipping points leading to major insurgent activities. In addition, realistic historical data was generated and correlated with simulation results to demonstrate how these models can portray real insurgent environments.

Keywords simulation · agent based model · insurgency · civil violence · graphics processing unit · distributed computing

1. Introduction

There has been increasing Department of Defense (DOD) interest in developing human, social, cultural, and behavioral (HSCB) models for understanding the complex human terrain in which many operations take place (Joint Warfighting Center, 2007). HSCB Models can be employed for simulating mission scenarios, determining optimal strategies for disrupting terrorist networks, or training and mission rehearsal (Zacharius et al., 2008). Models can also be used for understanding or predicting the effects of diplomatic, military, and economic courses of action on the attitudes and behaviors of a population in a region of interest.

For example, these models can provide a tool for assessing the response of both insurgent and civilian populations to the presence of friendly forces in a given military area. HSCB models for these types of social conflicts have been constructed in a variety of forms. These forms include statistical analyses of historical data (Gulden, 2002; Mannes, 2008), game theoretic models (Myerson, 1997; Goh et al., 2006), social network analyses (Krebis, 2002; Carley, 2004), and agent-based models (ABMs) (Epstein, 2002; Kuznar and Sedlmeyer, 2005). Although there are advantages and disadvantages to each approach (Parunak et al., 1998; Axtell, 2004), ABMs have seen considerable attention from the research and defense communities.

ABMs permit valuable insights into the workings of complex dynamic systems that may not be otherwise adequately captured by other methods. However, the utility of such models is implicitly tied to their ability to accurately and realistically portray the real world. Computational

restrictions may prohibit the ability to simulate an environment featuring large numbers of heterogeneous agents responding to complex influences. Recent modeling efforts (Chaturvedi et al., 2005; Holcombe et al., 2006; Parker, 2007) have made started using distributed clusters of computers to simulate large quantities of agents. However, this approach requires expensive, immobile resources with often limited availability. A novel approach has been introduced (D'Souza et al., 2007; Lysenko et al., 2007) that makes use of the distributed processing capability of a graphics processing unit (GPU).

In this report, we take advantage of this technique for simulating large-scale ABMs for a military regional insurgency application. The developed model simulates hundreds of thousands of agents that can exhibit a variety of cognitive and behavioral states as well as actions. The characteristics of this model are loosely based on the demographic and geophysical characteristics of Ramadi, Iraq, between April 2003 and August 2008.

This model introduces several innovations aimed at increasing realism versus real world human behavior. A hidden Markov process is used to model agent cognitive and behavioral states. Participation in insurgency and counterinsurgency efforts can be both active and passive. In addition, agents interact with the environment through constraints imposed by the incorporation of satellite imagery. Combining these novel behavioral elements with the high-performance, large-scale capabilities offered through the GPU leads to a model that can simulate realistic insurgent and counterinsurgent activities. Further, with this improved realism, the potential of such models is improved for operations planning and analysis needs..

2. Materials and Methods

2.1. Background and General Framework

ABMs employ a “bottom-up” approach (Epstein, 2006) to modeling, where heterogeneous entities (e.g. agents) are individually governed by personality- and cognition-dependent behavioral rule sets. Each agent responds to inputs from its local environment and some subset of other agents without centralized control. This approach to modeling offers an intuitive way of representing a large number of individual, distributed, and decentralized active objects, while making the complex human terrain explicitly more tractable. ABMs provide a flexible structure, whereby differing, complex, and nonlinear global behaviors can manifest through simple changes to an agent’s cognition, relationship with other agents, or ability to adapt through its observations. In addition, ABMs can more easily incorporate local (in space or time) characteristics in a system than a model based on aggregate statistical properties with centralized control.

Recently, Epstein (2002) introduced the idea of using ABMs for understanding and analyzing human behavior in a civil violence paradigm. This model employed two types of agents: an agent that can become either an active insurgent or quiet, and a police agent that can arrest nearby active insurgents. Epstein showed that by using even simple rules governing these agents’ behaviors, complex nonlinear behaviors could emerge. Although not explicitly coded into the model, Epstein was able to show instances of individual deceptive behavior, local (in space/time)

outbursts characteristic of mobs or riots (DeNardo, 1985), salami tactics of corruption, and tipping points (Kuran, 1989; Gladwell, 2002) leading to widespread insurgent activities. In this work, a second model was also introduced where there are two culturally-combative groups. In this extension, he was able to show emergent behaviors ranging from peaceful coexistence to ethnic cleansings, and how an active police force could result in the materialization of regional safe havens where one group was protected from the other.

Several other research efforts (Yiu et al. 2002; Taylor et al., 2004; Bulleit and Drewek, 2005; Kuznar and Sedlmeyer, 2005; Goh et al. 2006; Cui and Potok, 2007) have followed aimed at modeling similar insurgent environments. These models have introduced different motivations affecting agent actions, more complex movement behaviors, or capabilities of agents to “learn” from past experience. While these efforts are progressions toward more realistic simulations of social conflict, there are still limitations in these models.

For this report, we attempt to address some of these shortfalls through several advancements. We similarly build upon the Epstein civil violence model to incorporate more complex decision-making processes and behaviors of individual agents. In addition, the environment in which agents operate has geophysical characteristics previously unaccounted for in other insurgency models. Furthermore, the introduced model is capable of simulating very large numbers of agents within a very large regional environment by taking advantage of GPU distributed computing capabilities. The geophysical environment, agent demographics, and historical data are loosely connected to the social conflict environment in Ramadi, Iraq between April 2003 and August 2008. These novel extensions increase the viability of the model for portraying such a real world environment.

2.2. Agent-based insurgency model

For the created model, two different categories of actors are specified. One category makes up the general population of the region, and represents both those who are and are not actively rebellious. The other category represents the “peacekeeper” forces, which seek out and arrest actively rebellious agents. Agents making up the general population can dynamically and heterogeneously make decisions about whether they choose to participate in insurgency or counterinsurgency efforts. The peacekeeping force is modeled more simply with the constant objective to arrest active insurgents. Both categories of agents are capable of moving and performing actions against other agents in the system.

2.2.1. Initial agent states

As in Epstein (2002), an agent’s tendency to become an insurgent is governed by its feelings of grievance and its proclivity to take risks. We represent the agents’ personal decision-making process by a hidden Markov model (Schrodt, 2000; Weaver et al., 2001; Liu and Salvucci, 2001), whereby an agent has a hidden state and an observed (exposed) state (Figure 1). Markov processes have advantages for representing human decision-making in that they have a dynamic and stochastic structure, such that they are specifically designed to deal with the uncertainties common to human cognition and the influences of past behaviors (Schrodt, 2004).

For this model, the hidden state represents an agent’s perceived grievance, while the exposed state determines its behavior at any given time. At the beginning of a simulation, an agent’s

grievance is a function of two parameters, meant to simply encapsulate an agent's cognition. An agent's initial perceived hardship H_0 (e.g. physical, economic, social, or educational privation) is heterogeneously initialized as:

$$H_0 = U(\alpha_H, \beta_H), \quad (1)$$

where $U(\alpha, \beta)$ describes a random uniform distribution from α to β , such that $-1 \leq \alpha \leq \beta \leq 1$. Another crucial factor affecting grievance is the exogenous perceived legitimacy L_0 ($-1 \leq L_0 \leq 1$) of the regime or peacekeeping force (if some element of the peacekeeping force is of foreign origin). The initial grievance of an agent is:

$$G_0 = H_0(1 - L_0), \quad (2)$$

where L_0 is constant across all agents. From an agent's initial grievance, the hidden state decision rule can be defined as:

if $G_0 > T_D$, agent is *disgruntled* (D)
 if $T_S < G_0 < T_G$, agent is *neutral* (N)
 if $G_0 < T_S$, agent is *satisfied* (S)

where T_D and T_S are grievance thresholds, such that $0 \leq T_S < T_D \leq 1$.




Figure 1. Decision-making process of agents in the general population, represented as a hidden Markov process.

Each agent has a hidden state governed by its grievance, and an exposed state that determines its behavior, and governed by its hidden state and its willingness to take risks. The exposed state determines an agent's actions at each discrete time.

Given an agent's perceived grievance, an agent must decide if it wishes to act upon that grievance. It is therefore useful to define an agent's heterogeneous risk aversion R , which can be assumed (like H) to range from $U(\alpha_R, \beta_R)$. However, even those with a penchant for risk will be cautious if the circumstances of their local environment make action quite dangerous. An aggrieved agent will estimate the likelihood of getting arrested before determining whether they choose to actively rebel. Alternatively, a satisfied agent will choose whether to aid peacekeepers

based on the chance of getting attacked by an insurgent. Given an agent's local environment, limited by its vision V (defined as the radial Euclidean distance an agent can inspect), we can describe an agent's arrest (by a peacekeeper) and attack (by an active insurgent) probabilities based on the number and proximity of other agents nearby:

$$\begin{aligned} P_{i,\text{arrest}} &= 1 - e^{-k(PK_{n-1}/AI_{n-1})_V}, \\ P_{i,\text{attack}} &= 1 - e^{-k(AI_i/C_{il})_V}, \end{aligned} \quad (4)$$

where k is a constant set to ensure a plausible probability estimate. PK_{n-1} is a measure representing the proximity of peacekeeper agents; similarly, AI_{n-1} represents the proximity of active insurgents and C_{n-1} represents the rest of the population. These parameters are each computed for a given agent i with vision V . To ensure both the quantity and distance to the other agent types are considered, these terms are computed as:

$$\begin{aligned} PK_i &= N_{PK} \left[1 - \frac{1}{N_{PK}} \left(\sum_{j=1}^{N_{PK}} |\mathbf{p}_j - \mathbf{r}_i| \right) \right], \\ AI_i &= N_{AI} \left[1 - \frac{1}{N_{AI}} \left(\sum_{j=1}^{N_{AI}} |\mathbf{p}_j - \mathbf{r}_i| \right) \right], \\ C_i &= N_C \left[1 - \frac{1}{N_C} \left(\sum_{j=1}^{N_C} |\mathbf{p}_j - \mathbf{r}_i| \right) \right], \end{aligned} \quad (5)$$

where N_{PK} , N_{AI} , and N_C are the respective numbers of peacekeeper, active insurgent, and other general population members within an agent's vision. $\mathbf{r}_i$ is the current position vector of the given agent i , and $\mathbf{p}_j$ are the positions of the other types of agents within agent i 's vision.

The exposed (behavior) state decision is determined based on these four factors: an agent's hidden (grievance) state, risk aversion R , likelihood of getting accosted P , and a homogeneous risk activity threshold T_R (Table 1). As such, the root causes of agent behavior are delineated based on their sympathies toward insurgent or peacekeeper activities (grievance) and willingness to take action (net risk). A disgruntled agent identifying with insurgent activities must decide whether to act covertly or openly. Conversely, a satisfied agent must weigh the risks of assisting peacekeeper actions. A neutral agent may lean toward one group or another if safety concerns lend prudence to such leanings.

It should be noted that all of these agent attributes are abstractions. For example, the hardship could be qualitatively defined as some cumulative combination of various regional demographic features such as age, income, religion, ethnicity, tribe, education, occupation, or ideology (Chaturvedi et al., 2005). As a homogeneous and exogenous attribute, the perceived legitimacy can similarly take a variety of analogous global factors into effect. For this work, more focus was placed on the distinction in how these parameters are defined within the context of the model (e.g. homogeneously vs. heterogeneously across agents), and less what these parameters abstractly represent relative to real world statistics or demographics. If one were to try to

incorporate real world statistics into agent attributes, the semantic relationships between these attributes and these statistics would have to be well-defined.

Table 1. Agent exposed state decision, based on their hidden state (grievance), risk aversion R , probability of risk P , and risk threshold T_R .

Grievance State	Decision	Exposed State
Disgruntled (D)	$RP_{\text{arrest}} < T_R$	Active Insurgent (AI)
Disgruntled (D)	$RP_{\text{arrest}} \geq T_R$	Quiet Insurgent (QI)
Neutral (N)	$RP_{\text{arrest}} < RP_{\text{attack}}$	Quiet Insurgent (QI)
Neutral (N)	$RP_{\text{arrest}} \geq RP_{\text{attack}}$	Quiet Civilian (QC)
Satisfied (S)	$RP_{\text{attack}} < T_R$	Active Civilian (AC)
Satisfied (S)	$RP_{\text{attack}} \geq T_R$	Quiet Civilian (QC)

2.2.2. Agent cognitive dynamics

Over time, an agent can come across prosperity, or their estimation of perceived governmental activities can change. Therefore, we define terms for the dynamic change in an agent's grievance by:

$$G_n = G_{n-1} + \Delta H(1 - \Delta L) = G_{n-1} + \Delta G, \quad (6)$$

such that G_n is constrained between zero and one. G_{n-1} is a given agent's previous grievance, G_n is an agent's current grievance, and ΔH , ΔL , and ΔG are the respective changes in perceived hardship, legitimacy, and grievance. Like the initial hardship, the change in hardship is modeled as a stochastic process such that $\Delta H = U(\alpha_{\Delta H}, \beta_{\Delta H})$. This change will typically be much smaller than the initial hardship, but always between -1 and +1. The perceived legitimacy in the model is allowed to change either linearly in time or instantaneously triggered by some extreme or catastrophic event. The changes in perceived legitimacy were varied in different simulation runs to show how these changes can result in different behaviors.

Though risk aversion R in this model remains constant for each heterogeneous agent throughout the simulation, each agent reassesses its probability of danger based on changes to agents in its local neighborhood using Eqn (4). In addition, agent's movements and other actions will affect simulation dynamics, as will be seen in the next subsections.

2.2.3. Agent Movement

Like Epstein (2002), the constructed model consists of a two-dimensional lattice structure. Others have developed ABMs in a continuous space with agents represented by infinitesimal particles (Reynolds, 1999; Mitchell, 2008). There are advantages and disadvantages to both approaches, however, a lattice approach was chosen for this model so that movement decisions are represented through the appropriate level of granularity desired for a regional insurgency

application. The model involves a very fine grid (1184 x 611), where high resolution can still be achieved despite the lattice structure.

The lattice is backed by an image of a geospatial region. In this case, a satellite image was taken from Google Maps (www.maps.google.com) of the region around Ramadi, Iraq (Figure 2). A black and white mask of the city and surroundings was created such that the density of white pixels is higher in the more densely populated urban regions (Figure 3). Also in the mask, obstacles such as rivers were colored black. Agent spawning points were determined by (randomly) uniformly allowing agents to inhabit any white area of the lattice. Thus, agents will be initialized primarily in high density regions of the map (Figure 4), with agents allowed to occupy the same grid space.



Figure 2. Google maps image employed of Ramadi, Iraq and surrounding regions.

The Epstein (2002) model incorporated a relatively small (40x40) lattice with toroidal structure, such that agent movement that extended beyond a grid edge would appear on the opposite side of the lattice. This was necessary for that simulation, as the number of movement positions is severely limited by the smaller grid size. However, people in the real world can leave a region that is represented by a model as a closed system. Therefore, the model developed for this work allows agents to walk off grid edges as well as return; yet, they will not be visually displayed within the lattice boundaries and will not be able to “see” (in their vision) other agents also outside this boundary. Hence, agents can choose to leave the simulated region, for example, to flee from an oppressive environment. Nevertheless, agents do possess an impetus to remain within the grid region, as will be explained later in this section. In addition, a second image mask was created that defines available movement locations (Figure 5). Thus, agents are prevented from moving into grid locations containing obstacles such as rivers.



Figure 3. Black and white mask of the simulation region such that urban areas have a higher density of white pixels, while rural areas and obstacles such as rivers are primarily black.



Figure 4. Sample spawning distribution of agents, where agents are uniformly positioned in any white region of the image from the mask shown in Figure 3.

In ABMs, movement is typically defined by one of three methods. An agent's movement can be based purely on a random walk through available movement locations (Epstein, 2002). Alternatively, movement can be defined by velocity vectors in desired directions (Reynolds, 1999). Thirdly, an agent's movement can be based on the relative attractiveness of various candidate movement positions (Gill and Grieger, 2003).

The “random walk” approach may be practical for models where movement is considered less imperative than other simulation characteristics. The relative importance of directed movement is often determined by its sensitivity to key model outputs or based on the desired fidelity of the model to be constructed. The other two approaches rely on incorporating this directed movement, defined through attractive and repulsive forces. The vector-based approach typically assumes a continuous, particle-based model. This approach has the highest granularity in that it incorporates physical dynamic properties of the moving agents. However, for an insurgency model where discrete time is more likely to be measured in weeks rather than seconds, this level of sophistication in the movement dynamics may be computationally infeasible.



Figure 5. Black and white mask of the simulation region where obstacles such as rivers are not permissible movement locations (represented by black pixels).

Consequently, we employ movement rules based on the lattice-based distillation algorithm introduced by Gill and Grieger (2003). This approach computes the fitness of nearby lattice positions. For each agent position, the same vision used in cognition determines the radius of available movement locations. For each of these available lattice positions, a penalty function can be calculated based on each agent’s desire to seek out certain types of agents (i.e. agents of a certain exposed state) and flee from other types (i.e. agents of different exposed states). In addition, an impetus to move toward “home” (their initial spawning point) is also integrated into the movement penalty function, such that agents are always attracted to some extent to where they live. The complete penalty function for a given agent i and movement position $move$ is as follows:

$$f_{i,move} = w_{seek}D_{i,move,desired} - w_{flee}D_{i,move,undesired} + w_{home}D_{i,move,home}, \quad (7)$$

where the w parameters represent relative weights of influence on movement and the D parameters reflect distance measures to desired agents, undesired agents, and an agent’s home. The distances are fully described by:

$$\begin{aligned}
D_{i,move,desired} &= \sum_{j=1}^{N_{desired}} \left(\frac{|\mathbf{p}_{j,desired} - \mathbf{r}_{i,move}| - |\mathbf{p}_{j,desired} - \mathbf{r}_{i,stay}|}{|\mathbf{p}_{j,desired} - \mathbf{r}_{i,move}|} \right), \\
D_{i,move,undesired} &= \sum_{j=1}^{N_{undesired}} \left(\frac{|\mathbf{p}_{j,undesired} - \mathbf{r}_{i,move}| - |\mathbf{p}_{j,undesired} - \mathbf{r}_{i,stay}|}{|\mathbf{p}_{j,undesired} - \mathbf{r}_{i,stay}|} \right), \\
D_{i,move,home} &= \frac{|\mathbf{p}_{home} - \mathbf{r}_{i,move}| - |\mathbf{p}_{home} - \mathbf{r}_{i,stay}|}{|\mathbf{p}_{home} - \mathbf{r}_{i,stay}|},
\end{aligned} \tag{8}$$

such that there are N agents (within vision) of either desired or undesired states at lattice positions $\mathbf{p}_j$. Likewise, $\mathbf{p}_{home}$ is the lattice position of an agent's home (initial position at the beginning of the simulation). $\mathbf{r}_{i,move}$ is the candidate position an agent can move to, and $\mathbf{r}_{i,stay}$ is an agent's current position. Thus, each distance computation inside the sum is the relative change in proximity to other agents or to home. Summing these distance computations allows both the density and proximity of agents to factor into the penalty function. The penalty for staying in the same place will always be zero. A negative penalty represents a lattice position that is preferable based on this proximity to other agents. Conversely, a positive penalty is undesirable. It should be noted that an agent's assessment of available lattice positions is based on local knowledge of other agent positions, and a lattice point with a lower penalty may actually be worse because undesired agents are nearby but just outside an agent's current vision.

The agent types sought and fled from will depend on the agent state decisions described in Section 2.2.1. For this model, an agent of a given state will have an impetus to either seek or flee a certain other agent type, but not both. The implemented agent movement influences are shown in Table 2.

Table 2. Agent movement influences based on agent exposed state decision (agent type).

Agent Type	Desired Agents	Undesired Agents	Seek Home?
Active Insurgent (AI)	None	PK	Yes
Quiet Insurgent (QI)	AI, QI	None	Yes
Quiet Civilian (QC)	None	AI	Yes
Active Civilian (AC)	PK	None	Yes
Peacekeeper (PK)	PK, AI	None	Yes

To introduce a random element into agent movements, a binary tournament selection algorithm (Miller et al, 1995) is implemented. This approach is applied regularly in optimization problems to avoid getting trapped in local optimal conditions, when a better global condition is achievable (Mitchell, 1998). Furthermore, allowing a stochastic element permits agents to act more like realistic entities, where motivations may not be fully captured by the abstracted parameters represented in the model. For a binary tournament, two of the possible movement positions are

randomly and uniformly selected. The agent then moves to the position of the two selected that “wins the tournament” by having the lower penalty function.

2.2.4. Agent Actions

Agents from the general population can act in four possible ways, completely determined by their exposed states (Figure 1). Active insurgents can attack other agents, either civilians or peacekeepers. Whom they attack is decided based on chance and opportunity. An active insurgent decides to attack a peacekeeper or civilian each with a probability of 0.5. If no agents of that type are within their vision, they will do nothing. If their vision finds only one agent of that type, they will attack that agent. If two or more agents are present, a binary tournament selection approach is employed as with movement (Section 2.2.3), where the closer agent is more likely to be attacked.

Quiet insurgents will aid active insurgents by providing an abstraction of “intelligence” information. For every quiet insurgent, an active insurgent’s vision is increased during the current time by an increment ΔA , where:

$$V_{\text{new}} = V_{\text{old}} + \sqrt{\Delta A / \pi} , \quad (9)$$

such that the new vision radius V_{new} is incremented proportional to the two-dimensional area described by this vision. Active civilians will similarly inform peacekeepers to increase their vision. Meanwhile, passive civilians will remain neutral and do nothing, although they are still susceptible to attack by active insurgents.

Peacekeeper agents will arrest active insurgents within their vision, and keep them for the jail time specified by the parameter J . Their decision on whom to arrest also uses a binary tournament selection algorithm where closer active insurgents are preferred.

2.3. Programming with a GPU

Compute Unified Device Architecture (CUDA) is NVIDIA Corporation’s software development model for General Purpose Programming on Graphics Processing Units (GPGPU) (NVIDIA Corporation, 2008). With CUDA, general purpose applications written in C or C++ may offload computationally intensive tasks to the GPU, exploiting the significant parallel processing capabilities that these low-cost devices provide.

A typical mid-range GPU currently has between 14 and 16 multiprocessors, with each multiprocessor comprised of 8 streaming processor cores. In Flynn’s taxonomy of processor architectures (Wilkinson and Allen, 2005), each multiprocessor has a Single Instruction, Multiple Data architecture as opposed to multi-core central processing units (CPUs) which have a Multiple Instruction, Multiple Data architecture. Thus, CUDA allows the GPU to handle problems that involve large amounts of data that can be partitioned into many separate chunks, and then processed using identical instructions. Common examples of such problems include large matrix transforms, data mining algorithms, and simulation of particle systems. In CUDA, all streaming processor cores of a given multiprocessor must execute the same instruction on separate chunks of data; otherwise the instructions must be serialized for each processor.

ABMs share much in common with simulations of particle systems. Each agent can be thought of as a particle that interacts with other agents according to a predefined set of rules. If a GPGPU solution could be implemented, it would allow the simulation to scale upward from only a few hundred agents, as is common in open-source ABM frameworks such as Repast (North et al., 2005), Swarm (Minar et al., 1996), or MASON (Luke et al., 2004), to tens or hundreds of thousands of agents.

To achieve this, we implemented a number of non-trivial parallel algorithms. An example of this is the vision (V) algorithm (Section 2.2.1) that allows agents to sense the presence of other agents in their environment. A brute force approach to this algorithm would involve each agent examining every other agent in the simulation and then testing whether the other agent was close enough to detect: an $O(n^2)$ solution. Our implementation, which is based on a number of techniques found in (LeGrand, 2008), results in an $O(n \log n)$ solution, allowing a significant reduction in processing time.

The most obvious solution would be to create an array of length equal to the number of lattice points in the model, where each index contains a pointer to the agent located in the corresponding lattice point. Assuming threads are allocated on a per-agent basis, each thread need only reference indexes relative to its agent's position. However, since this model allows varying numbers of agents to occupy a single lattice point at any time, another approach is required, as shown in Figure 6.




Figure 6. Data structures implementing the parallel agent vision algorithm.

A Large Scale, High Resolution Agent-Based Insurgency Model

This example shows a 4x4 model lattice occupied by a total of eight agents. Each agent has a unique Agent ID, numbering 1 to n , and each lattice point, or cell, is given a unique integer value based on its x and y coordinates according to the function:

$$f(x, y) = x + (y * \text{lattice width}). \quad (10)$$

To allow each thread to properly reference nearby agents, two arrays are created: a Cell ID array of length equal to the number of agents in the model and a Cell Index array of length equal to the number of lattice points in the model.

Indexes in the Cell Index array contain an index into the Cell ID array and an offset describing how many agents may be indexed from that position. In the example from Figure 6, lattice point 6 which is located in column 2 of row 1, is occupied by two agents whose IDs are 2 and 7. Index 6 of the Cell Index array declares that lattice point 6 has two agents, and that these may be found beginning at index 4 of the Cell ID array.

During execution, each thread associated with a given agent traverses the Cell Index Array to examine nearby lattice points for the presence of other agents. If a presence is indicated in the Cell Index Array, then a thread may examine the agent(s) present there via the Cell ID Array.

3. Results

Simulation results will be presented by showing some interesting examples of manifest behaviors the model is capable of capturing. These examples will be presented to show the breadth of emergent insurgency characteristics that can be demonstrated and their associations to real life. The advantage of using the GPU approach to model development can be illustrated by the massive agent population and lattice size. Six different simulation runs will be presented (see Table 3 for simulation parameters for each run).

3.1. Run A: Localized Outbursts

Many interesting local spatial system behaviors can be captured in the agent-based insurgency model. For example, in zones with a limited peacekeeper presence, high densities of active insurgents can concentrate. This behavior can be visualized in Figures 7 and 8. Although active insurgents are not implicitly motivated to move toward other active insurgents, the presence of other actives in the nearby vicinity reduces the risk of arrest from peacekeepers. To be the first active insurgent, one must have a high grievance and low risk aversion. However, the presence of active insurgents provides a perpetual feedback mechanism for bolstering other aggrieved agents to also become active within their vicinity. Quiet insurgents in this region are more likely to become active, and civilians are likely to flee to safer areas.

Another interesting property in this depiction is the tendency for these localized assemblies to produce fronts near geographically isolated borders, such as along river banks with limited inland access. These regions have reduced access for peacekeeper agents to enter, thus the density of peacekeepers in their vicinity will tend to be inherently low. Thus, the self-organizing active insurgent zones are more likely to be formed in these environments with low risk of arrest. Likewise, peacekeeper agents searching out active insurgents will have only limited access

A Large Scale, High Resolution Agent-Based Insurgency Model

points from which to travel into these local zones, bottlenecking infiltration efforts. Thus, the numbers and locations from which peacekeepers can disrupt the formation of these active assemblies are limited. What emerges is “siege”-like behaviors, whereby a small number of peacekeepers gather outside these isolated locales while active insurgents cluster tightly inside these bordered regions.

Table 3. Model parameters for the various simulation cases used in Sections 3.1 – 3.6.

Parameter	Run A	Run B	Run C	Run D	Run E	Run F
Agent Population Size, N	100,000	100,000	100,000	100,000	100,000	100,000
Initial Peacekeeper Density, N_{PK}/N	10%	10%	1: 10% 2: 20% 3: 10%	10%	10%	2%
Peacekeeper Vision, V_{PK}	5.0	3.5	1: 3.5 2: 3.5 3: 7.0	3.5	3.5	3.5
Other Agent Vision, V	5.0	3.5	3.5	3.5	3.5	3.5
Perceived Legitimacy, L_0	0.15	0.87	0.80	1: 0.90 2: 0.89	0.25	3.5
Min & Max Initial Hardship, α_H and β_H	[0,1]	[0,1]	[0,1]	[0,1]	[0,1]	[0,1]
Satisfied Threshold, T_S	0.25	0	0	0	0.25	0
Disgruntled Threshold, T_D	0.5	0.1	0.1	0.1	0.5	0.1
Min & Max Risk Aversion, α_R and β_R	[0,1]	[0,1]	[0,1]	[0,1]	[0,1]	[0,1]
Risk Threshold, T_R	0.01	0.01	0.01	0.01	0.01	0.01
Min & Max Hardship Change, α_{AH} and β_{AH}	[-0.01, 0.01]	[-0.01, 0.01]	[-0.01, 0.01]	[-0.01, 0.01]	[-0.1, 0.1]	[-0.01, 0.01]
Linear Legitimacy Change, ΔL	N/A	N/A	N/A	1: -0.0005 2: N/A	N/A	-0.0001: for first 200 time steps 0.002439: for last 82 steps
Instant Legitimacy Change, $\Delta L(n)$	N/A	N/A	N/A	1: N/A 2: -0.2	N/A	N/A
Time for Instant Legitimacy Change, n	N/A	N/A	N/A	1: N/A 2: 400	N/A	N/A
Seek/Flee Movement Weight, w_{seek} and w_{flee}	1.0	1.0	1.0	1.0	1.0	1.0
Home Movement Weight, w_{home}	0	0	0	1.0	0	1.0
Vision Increase Per Informing Agent, ΔA	0.01	0	0	0.01	0	0.01
Jail Term, J	1	30	30	∞	1	1




Figure 7. Regions with localized outbursts of active insurgent activities, indicated by red pixels (blue pixels are peacekeepers and green pixels are other members of the general population). Insert box corresponds to Figure 8.



Figure 8. Insert from Figure 7: Regions with localized outbursts of active insurgent activities, indicated by red pixels (blue pixels are peacekeepers and green pixels are other members of the general population).

3.2. Run B: Punctuated Equilibrium

Figure 9 shows the number of active insurgents in time for a simulation run that spans a large number of time steps. The system tends to dwell predominately in a state of relative stability, where active insurgents typically number below about 50. However, there are sharp spikes of larger-scale insurgent activity throughout the simulated duration. This dynamic is representative of punctuated equilibrium, where the described periods of relative stability are punctuated by episodic outbursts of rebellious activity (Young, 1998). This corresponds with Epstein's work (2002), who first showed how agent-based insurgency models could capture this type of nonlinear behavior.




Figure 9. Population dynamics often characterized by punctuated equilibrium.

3.3. Run C: Effect of Peacekeepers

Many noteworthy population dynamics can be generated through modifications to various model parameters. For example, if a model were developed for mission operations and planning, one may modify the locations and behaviors of peacekeepers to examine the effect of different strategies on the overall system behaviors. In this set of simulation runs, we compared simulations with peacekeeper agents representing 10% and 20% of the total population, respectively. Figure 10 shows the number of agents of various exposed states over time, as well as the number of active insurgents arrested by the peacekeepers. The number of quiet insurgents (Figure 10(a)) stays constant through both simulations, but the numbers are much lower with a higher percentage of peacekeepers. In this case, neutrally aggrieved agents weigh the risks of passively assisting insurgents or staying neutral, and tend to stay neutral when the density of peacekeepers is high.

Likewise, highly aggrieved agents are less likely to take the risk of becoming active when the probability of getting arrested is higher (Figure 10(b)), which in turn, prevents the perpetual feedback of described earlier, where the presence of actives leads to more actives. Despite the larger peacekeeping force, the number of arrests stays fairly similar because less members of the general population are willing to become active (Figure 10(c)).

An alternative maneuver might be to supply the peacekeepers with better technology and resources for gathering intelligence information. This intelligence advantage is modeled abstractly as a relative increase in the peacekeeper's vision V . In this case, the number of aggrieved agents and agents willing to aid peacekeepers is the same, as it is based on local sizes of the peacekeeper force (10% peacekeeper density). However, the number of active insurgents remains low. Because of their larger vision, peacekeeper agents can more quickly find and stamp down individual insurgent sparks before they escalate into larger outbursts.




Figure 10. Population dynamics when the peacekeeper parameters are different, where (a) number of quiet insurgents in time, (b) number of active insurgents in time, and (c) number of arrested insurgents in time.

3.4. Run D: Salami Tactics of Corruption

In Figure 11, the perceived legitimacy of the general population was decreased in two different ways. First, the initial legitimacy was decremented linearly through the simulation from $L_0 = 0.9$ to $L = 0.4$ at the end of the simulation. Second, the legitimacy was left at a constant $L_0 = 0.89$,

and then instantaneously reduced to $L = 0.7$ after 400 time steps. It is quickly notable that the number of quiet insurgents reflects these changes (Figure 11(a)), as the grievance (assuming a fairly constant hardship) will be proportional to the legitimacy. Conversely, the number of civilians has the inverse relationship to the number of quiet insurgents (Figure 11(c)).




Figure 11. Population dynamics when there is a gradual or instantaneous perceived legitimacy decline, where (a) number of quiet insurgents in time, (b) number of active insurgents in time, (c) number of civilians in time, and (d) number of arrested insurgents in time.

When the legitimacy declines very slowly, the number of active insurgents at any one time remains relatively small and steady (beyond an initial transient). Agent grievances increase very slowly, so active insurgents likewise appear incrementally. Peacekeeper agents are able to arrest these new actives as they appear, thereby preventing clusters of active insurgents from forming (in contrast to what was shown in Figures 7 and 8). The arrested population rises smoothly over time as actives appear, such that there are no significant spikes in the number of actives at any given time.

Alternatively, when the population feels a sudden reduction in the perceived legitimacy of those in power, there is an immediate response in the number of quiet and active insurgents in the population. Consequently, pockets with dense active insurgent populations appear quickly, and cannot be stamped down efficiently by the peacekeeping force. There is a steep rise in the

number of actives and the number of arrests by peacekeepers until stability can be regained much later in the simulation. Even though the absolute legitimacy reduction is smaller in the second scenario, the insurgency event is much greater.

This result corroborates the notion that a corrupt regime is more likely to be accepted by the majority if it chips away at the legitimacy of a government. These types of “salami tactics” were employed by Adolf Hitler’s Nazi Party in the early 1930s and by Stalinist Mátyás Rákosi of the Hungarian Communist Party in the late 1940s. Alternatively, one can surmise how dramatic “triggering events” such as assassinations (Kuran, 1989; Epstein, 2002) can result in a flood of civil violence.

3.5. Run E: Effect of Geography

One element that is often left out of other ABMs is the effect of geography. We have already showed how movement constraints can lead to more localized outbursts at distinct locations (Section 3.1). For this model, the initial agent positions and allowable movement locations are determined by masks as explained in Section 2.2.3. To examine what happens when different spawning and movement masks are utilized, we conducted a scenario where parameters were unvaried but different geographical masks were used. Figure 12 shows simulation runs executed with the same parameters but different agent spawning location and movement masks. Figure 12(a) displays a screen capture of the simulation run using the masks shown previously, created from the image of Ramadi. Figures 12(b) and (c) utilize a spawning mask from a different city (Kirkuk), where 12(b) strictly constrains agent movements and 12(c) allows very free movement. The numbers, concentrations, and locations of active insurgents during a simulation can be greatly affected by the masks that are chosen. Here, when agents are freer to move to nearly any lattice position, the amount of insurgent activity is much greater (Figure 13). This result occurs because quiet and active insurgents are more able to congregate together within their visions. In this simulation, when agents are forced to separate, either through geography or through counterinsurgency efforts, the strength of the insurgency can be reduced.

3.6. Run F: Matching Historical Data

To demonstrate the capabilities of this model, a realistic scenario was envisaged based on the events in Ramadi, Iraq (Figure 2) in the Al Anbar province, between April 2003 and August 2008. Ramadi is considered to be the southwest point of what is termed the “Sunni Triangle,” a region in Iraq occupied largely by Sunni Muslims and deemed to be an early focal point of the resistance to United States (US) forces during the War in Iraq. Unclassified documentation puts US operations in this area starting around March of 2003. From this time until the fall of 2006, the violence continued to escalate in this region, as the general population built an ever-increasing grievance against the US occupation.

During early 2007, US forces began to have success stamping down the rebellion, a reversal primarily attributed to an event called the “Anbar Awakening.” In late August 2006, insurgents shifted their tactics, and attacked an Iraqi Police Station. They murdered the Sheikh of the tribe primarily responsible for manning this station, and then violated Arab custom by leaving the body where it could not be found for days. This dramatic “triggering event” enraged rather than intimidated cooperative tribal leaders. These leaders met with US forces and formed the Al Anbar Awakening movement. This movement spurred an increase in the number of Iraqi

Security Forces recruits, while insurgent attacks dropped dramatically and cooperation with Coalition forces spread throughout the province. These events represented a prominent success story for the Iraq War counterterrorist campaign (Fumento, 2006; Smith and MacFarland, 2008).




Figure 12. Simulation screen capture at time step 250 using the same parameters but different geographical elements, where (a) Ramadi, Iraq spawning mask and movement mask, (b) Kirkuk spawning mask and movement restrictive movement mask, and (c) Kirkuk spawning mask and unrestrictive movement mask. Red pixels are active insurgents, blue pixels are peacekeepers, and green pixels are other members of the general population.

It is easy to ascertain that the complexity of real historical events make trying to fit a model to these events quite challenging. Although the simulation run described here is still an abstraction of the real historical event, it can be shown that even face validated model parameters can recreate a scenario similar to what may happen in the real world. Because the Ramadi region is fairly homogeneous in terms of religion and tribe, we did not create a second category of actor to represent a different demographic of the general population. In a region populated by both Sunni

and Shi'a Muslims, or multiple tribes, it may be more important to distinctly model differing or opposing categories of agents.




Figure 13. Population dynamics when the geographies are varied, represented by masks that determine initial agent positions and allowable movement locations, (a) number of quiet insurgents in time, (b) number of active insurgents in time, and (c) number of agents that leave the lattice during the simulation. Masks 1, 2, and 3 correspond to the geographies represented by Figures 11(a), (b), and (c), respectively.

The events occurring in Ramadi were simply abstracted to just changes in the legitimacy parameter for this model. The effect of the initial presence of US forces in the region was modeled as a gradual decrease in the perceived legitimacy of the government, as dissention toward these US forces escalated. However, the Anbar Awakening represented a significant shift in the perceived legitimacy of Al Qaeda. To model this event, the same legitimacy parameter was increased, representing the increased favor garnered by Coalition forces and Iraqi Security forces as the insurgency lost esteem.

Increased complexity could of course be achieved to better match the events of this region and time period. For example, the locations and sizes of US and Coalition forces could be historically retrieved and integrated into the simulation. Other demographic and detailed event-driven aspects could also be incorporated. However, to make the model results described here fairly generic, these types of detailed customizations were not made.

To validate the model simulation results, a representative time series of significant insurgent acts was created. Although this “historical data” does not correspond with real historical data, it matches both the anecdotal account of the events that transpired, and is proportional to open-source significant event data corresponding to the region. The representation of historical data that was created counts weekly insurgent acts from April 2003 through August 2008. The trends in the data show a gradual increase in the number of uprisings until the Anbar Awakening, at which point the attacks against US and Coalition forces began to taper off. To match these trends in the simulation, agent behavioral parameters were empirically “tuned” under some constraints. The result can be seen in Figure 14. Although the data used here is partially fabricated, the application toward realistic data is evident. Given constraints imposed by information such as demographics and intelligence data, an agent-based model may be derived to match global behaviors exhibited in the real system, such that it can be employed as a useful tool for analysts.




Figure 14. Example empirical match of simulation data with an example set of historical data: matching the number of insurgent acts in time.

4. Discussion

In this work, an ABM was detailed, capable of simulating large numbers of agents in a complex insurgency environment. The key contributions to the body of knowledge include the use of a GPU for distributed programming in an agent-based insurgency model, to dramatically improve performance capability. In addition, human cognition was modeled based on a hidden Markov process, where agent grievances represent the hidden state, and agent behaviors represent the exposed (observed) state. Agents were able to make decisions to actively or passively participate in insurgency or counterinsurgency activities. These agents also possess directed movement based on their exposed state and nearby environment, including constraints. The development of

A Large Scale, High Resolution Agent-Based Insurgency Model

these attributes for an ABM in an insurgency application signifies a movement toward more realistic simulations that can be employed for operations planning and analysis.

As demonstrations of the capability of this model, several example simulation results were shown. Several representative emergent behaviors were achieved, such as localized outbursts, states of punctuated equilibrium, tipping points for massive insurgent activities, and slow corruption leading to passivity of populations. The utility of the model was demonstrated by showing how simulation runs can generate different behaviors based on varied counterinsurgency efforts, and how geography can also have an impact on insurgent activities. In addition, realistic “historical data” was generated by examining historical accounts and limited open-source data corresponding to Ramadi, Iraq between April 2003 and August 2008. This data was then compared to outputs from the ABM simulation. By empirically tuning the micro-level behavioral parameters of agents, simulations were achievable that match the characteristics of the existing “historical data.”

Although the model developed for this work can achieve both fairly realistic simulations and a wide breadth of behaviors, the task of modeling and/or predicting human behavior is still a challenging and often uncertain proposition. There is a need for continued development of HSCB models, such that improvements can continue to be made in regard to producing validated models that can realistically portraying the real world.

Immediate attention should be given to a more rigorous validation of the model, to include work with more detailed historical data. The large number of tunable parameters in the model as well as the complexity of their interactions may make traditional validation techniques involving design of experiments difficult to accomplish. Therefore, techniques from the field of evolutionary computation, such as genetic algorithms (Mitchell, 1998) may be useful. Data could also be integrated by constraining agent attributes based on demographic or intelligence information. Such a model could be validated both heuristically, by updating agent characteristics at the front end, and statistically, by matching results to actual global behavioral data at the output.

Another possible improvement could be made by incorporating the model into a larger federation of models by means of the High Level Architecture (HLA) or by exposing the model as a service in a Service Oriented Architecture (SOA). One possible SOA platform for integration is aXiom, a system that is currently under development at the Space and Naval Warfare Systems Center Atlantic. Integration into these platforms would also require research into a semantic approach for connecting historical data and/or intelligence information with the model federation such that it can be updated autonomously.

Finally, the use of GPUs to provide high performance computing power also provides the opportunity to exploit their graphics processing capabilities for Graphical User Interface (GUI) improvements. The current 2D graphical representation of the model does not fully capture the complex interactions that occur between agents; therefore there are still a number of emergent properties that cannot be interactively visualized. The incorporation of three-dimensional graphical elements would allow analysts to perceive much greater detail in the complex web of social interactions among agents.

Acknowledgements

The authors acknowledge support and funding through the Space and Naval Warfare Systems Center Atlantic Innovation Program. This work was performed under the auspices of the Advanced Technology Transition Branch (Code 56150), Information Warfare Exploitation Systems Engineering Division, Intelligence Surveillance and Reconnaissance / Information Operations, Space and Naval Warfare Systems Center Atlantic.

References

Axtell, RL, *Why Agents? On the Varied Motivations for Agent Computing in the Social Sciences* (Working Paper No. 17). Washington, DC: Brookings Institution Center on Social and Economic Dynamics, 2004.

Bulleit, WM, Drewek, MW. An agent-based model of terrorist activity. *Proc. North American Association for Computational Social and Organizational Science (NAACSOS)*. Notre Dame, Indiana, June 26-28, 2005.

Carley, K. Estimating vulnerabilities in large covert networks using multi-level data. *Proc. 2004 Int'l Symposium on Command and Control Research and Technology*. San Diego, CA, June, 2004.

Chaturvedi, AR, Dolk, D, Chaturvedi, R, Mulpuri, M, Lengacher, D, Mellema, S, Poddar, P, Foong, C, Armstrong, B. Understanding insurgency by using agent-based computational experimentation: case study of Indonesia. *Proc. Agent 2005 Conference*. Argonne National Laboratory, Argonne, IL, 2005.

Cui, X, Potok, TE. A particle swarm social model for multi-agent based insurgency warfare simulation. *Proc. 5th Int'l Conf. on Software Engineering Research, Management, and Applications*: 177-183, 2007.

DeNardo, J. *Power in Numbers: The Political Strategy of Protest and Rebellion*. Princeton, NJ: Princeton University Press, 1985.

D'Souza, RM, Lysenko, M, Rahmani, K. Sugarscape on steroids: simulating over a million agents at interactive rates, *Proc. Agent 2007 Conference*, Argonne National Laboratory, Argonne, IL, 2007.

Epstein, JM. Modeling civil violence: an agent-based computational approach. *Nat'l Acad Sci Colloquium* 99(3): 7243-7250, 2002.

Epstein, JM. *Generative Social Science: Studies in Agent-Based Computational Modeling*. Princeton, NJ: Princeton University Press, 2006.

Fumento, M. Return to Ramadi. *The Weekly Standard* 12(11), November 27, 2006.

Gill, AW, Grieger, D. *Validation of Agent Based Distillation Algorithms* (Technical Report DTSO-TN-0476). Australia: Defense Science and Technology Organization, Land Operations Division, 2003.

Gladwell, M. *The Tipping Point: How Little Things Can Make A Big Difference*. New York, NY: Black Bay Books, 2002.

Goh, CK, Quek, HY, Tan, KC, Abbass, HA. Modeling civil violence: an evolutionary multi-agent, game theoretic approach. *Proc. IEEE Congress on Evolutionary Computation*. Vancouver, BC, Canada, July 16-21, 2006.

Gulden, TR. Spatial and temporal patterns in civil violence: Guatemala, 1977-1986. *Politics and Life Sciences* 21(1): 26-36, 2002.

Holcombe, M, Coakley, S, Smallwood, R. A general framework for agent-based modelling of complex systems. *Proc. European Conf. on Complex Systems*. Paris, France, November 14-18, 2005.

The Joint Warfighting Center. An effects-based approach: refining how we think about operations. *Joint Forces Quarterly* 44, 2007.

Krebis, VE. Mapping networks of terrorist cells. *Connections* 24(3): 43-52, 2002.

Kuran, T. *Sparks and prairie fires. A theory of unanticipated political revolution*. Public Choice 61: 41-74, 1989.

Kuznar, LA, Sedlmeyer, R. Collective violence in Darfur: an agent-based model of pastoral nomad/sedentary peasant interaction. *Mathematical Anthropology and Cultural Theory: An International Journal* 1(4), 2005.

LeGrande, S. Broad-phase collision detection with CUDA. In: *GPU Gems 3*, edited by Nguyen, H. Upper Saddle River, NJ: Addison-Wesley, 2008, p. 697-721.

Liu, A, Salvucci, D. Modeling and prediction of human driver behavior. *Proc. 9th Int'l Conf. on Human-Computer Interaction*. New Orleans, LA, August, 2001.

Luke, S, Cioffi-Revilla, C, Panait, L, Sullivan, K. MASON: a new multi-agent simulation toolkit. *Proc. 2004 SwarmFest Workshop*. 2004.

Lysenko, M, D'Souza, RM. A framework for megascale agent-based model simulations on the GPU, *Journal of Artificial Societies and Social Simulation*. accepted, in press [2008].

Mannes, A, Michael, M, Pate, A, Sliva, A, Subrahmanian, VS, Wilkenfield, J. Stochastic opponent modeling agents: a case study with Hezbollah, *Proc. Social Computing, Behavior Modeling, and Prediction*. Phoenix, AZ, April 1-2, 2008.

Miller, BL, Goldberg, DE. Genetic algorithms, tournament selection, and the effects of noise. *Complex Systems* 9(3): 193-212, 1995.

Minar, N, Burkhart, R, Langton, C, Askenazi, M. *The Swarm Simulation System: A Toolkit for Building Multi-Agent Simulations* (Working Paper 96-06-042), Santa Fe Institute, Santa Fe, NM, 1996.

Mitchell, B. Symbiotic adaptive multisimulation. *Master's Thesis in Computer Science*. Auburn University, 2008.

Mitchell, M. *An Introduction to Genetic Algorithms*. Boston, MA: The MIT Press, 1998.

Myerson, RB. *Game Theory: Analysis of Conflict*. Boston, MA: Harvard University Press, 1997.

North, MJ, Howe, TR, Collier, NT, Vos, RJ. The Repast Symphony runtime system, *Proc. Agent 2005 Conference*. Argonne National Laboratory, Argonne, IL, October, 2005.

NVIDIA Corporation. NVIDIA CUDA Programming Guide 2.0 [Online]. NVIDIA Corporation. http://www.nvidia.com/object/cuda_develop.html [15 Oct. 2008].

Parker, J. *A Flexible, Large-Scale, Distributed, Agent-Based Epidemic Model* (Working Paper No. 52). Washington, DC: Brookings Institution Center on Social and Economic Dynamics, December, 2007.

Parunak, HVD, Savit, R, Riolo, RL. Agent-based modeling vs. equation-based modeling: a case study and users' guide. *Proc. Multi-Agent Systems and Agent-Based Simulations (MABS '98)*: 10-25, 1998.

Reynolds, CW. Steering behaviors for autonomous characters. *Proc. Game Developers Conference*: 763-782, 1999.

Schrodt, PA. Forecasting conflict in the Balkans using hidden Markov models. *Proc. Annual Meeting of the American Political Science Association*, Washington, DC, August, 2000.

Schrodt, PA. *Patterns, Rules and Learning: Computational Models of International Behavior*. Vinland, KS: Parus.

Smith, N, MacFarland, S. Anbar awakens: the tipping point. *Military Review*, March/April, 2008.

Taylor, G, Frederiksen, R, Vane III, RR, Waltz, E. Agent-based simulation of geo-political conflict. *Proc. 16th Conf. Innovative Applications of Artificial Intelligence (IAAI)*. San Jose, CA, 2004.

Weaver, R, Silverman, BG, Shin, H, Dubois, R. Modeling and simulating terrorist decision-making: a “performance moderator function” approach to generating virtual opponents. *Proc. 10th Conf. on Computer Generated Forces and Behavioral Representation, SISO*. May, 2001.

Wilkinson, B., Allen, M. *Parallel Programming: Techniques and Applications Using Networked Workstations and Parallel Computers*. Upper Saddle River, NJ: Pearson Prentice Hall, 2005.

Yiu, SY, Gill, A, Shi, P. Investigating strategies for managing civil violence using the MANA agent based distillation. *Proc. Land Warfare Conf.* Brisbane, Australia, 2002.

Young, HP. *Individual Strategy and Social Structure*. Princeton, NJ: Princeton University Press, 1998.

Zacharius, GL, MacMillan, J, Van Hemel, SB, eds. *Behavioral Modeling and Simulation: From Individuals to Societies*. Washington, DC: National Academies Press, 2008.