



NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

**ANALYSIS OF COGNITIVE ARCHITECTURE
IN THE CULTURAL GEOGRAPHY MODEL**

by

Chin Chuan Ong

September 2012

Thesis Advisor:
Second Reader:

Christian J. Darken
Jason Caldwell

**This thesis was performed at the MOVES Institute
Approved for public release; distribution is unlimited**

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, 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, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE September 2012	3. REPORT TYPE AND DATES COVERED Master's Thesis	
4. TITLE AND SUBTITLE Analysis of Cognitive Architecture in the Cultural Geography Model			5. FUNDING NUMBERS	
6. AUTHOR(S) Chin Chuan Ong				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government. IRB Protocol number ____N/A____.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited			12b. DISTRIBUTION CODE A	
13. ABSTRACT (maximum 200 words) The Cultural Geography (CG) Model is a multi-agent discrete event simulation developed by TRAC-Monterey. It provides a framework to study the effects of operations in Irregular Warfare, by modeling behavior and interactions of populations. The model is based on social science theories; in particular, agent decision-making algorithms are built on Exploration Learning (EL) and Recognition-Primed Decision (RPD), and trust between entities is modeled to increase realism of interactions. This study analyzed the effects of these components on behavior and scenario outcome. It aimed to identify potential approaches for simplification of the model, and improve traceability and understanding of entity actions. The effect of using EL/RPD with/without trust was tested in basic stand-alone scenarios to assess its impact in isolation on entities' perception of civil security. Further testing also investigated the influence on entity behavior in the context of obtaining resources from infrastructure nodes. The findings indicated that choice of decision-making methods did not significantly change scenario outcome, but variance across replications was greater when both EL and RPD were used. Trust was found to delay the rate of change in population stance due to interactions, but did not affect overall outcome if given sufficient time to reach steady state.				
14. SUBJECT TERMS Cultural Geography, Agent-Based Modeling, Human Behavior Modeling, Social Simulations, Reinforcement Learning, Recognition Primed Decision, Trust			15. NUMBER OF PAGES 79	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UU	

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

**ANALYSIS OF COGNITIVE ARCHITECTURE
IN THE CULTURAL GEOGRAPHY MODEL**

Chin Chuan Ong
Major, Singapore Armed Forces
B.Sc., University of Illinois at Urbana-Champaign, 2005

Submitted in partial fulfillment of the
requirements for the degree of

**MASTER OF SCIENCE IN
MODELING, VIRTUAL ENVIRONMENTS, AND SIMULATION (MOVES)**

from the

**NAVAL POSTGRADUATE SCHOOL
September 2012**

Author: Chin Chuan Ong

Approved by: Christian J. Darken
Thesis Advisor

Jason Caldwell
Second Reader

Christian J. Darken
Chair, MOVES Academic Committee

Peter J. Denning
Chair, Department of Computer Science

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

The Cultural Geography (CG) Model is a multi-agent discrete event simulation developed by TRAC-Monterey. It provides a framework to study the effects of operations in Irregular Warfare, by modeling behavior and interactions of populations. The model is based on social science theories; in particular, agent decision-making algorithms are built on Exploration Learning (EL) and Recognition-Primed Decision (RPD), and trust between entities is modeled to increase realism of interactions. This study analyzed the effects of these components on behavior and scenario outcome. It aimed to identify potential approaches for simplification of the model, and improve traceability and understanding of entity actions. The effect of using EL/RPD with/without trust was tested in basic stand-alone scenarios to assess its impact in isolation on entities' perception of civil security. Further testing also investigated the influence on entity behavior in the context of obtaining resources from infrastructure nodes. The findings indicated that choice of decision-making methods did not significantly change scenario outcome, but variance across replications was greater when both EL and RPD were used. Trust was found to delay the rate of change in population stance due to interactions, but did not affect overall outcome if given sufficient time to reach steady state.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I.	INTRODUCTION.....	1
A.	BACKGROUND	1
B.	PROBLEM STATEMENT.....	3
C.	OBJECTIVES.....	4
D.	METHODOLOGY	4
II.	OVERVIEW OF THE CULTURAL GEOGRAPHY MODEL	7
A.	DEVELOPMENT	7
B.	UNDERLYING CONCEPTS AND THEORIES	8
1.	Theory of Planned Behavior	8
2.	Narrative Paradigm.....	10
3.	Homophily	11
4.	Decision Making and Learning	12
a.	Reinforcement Learning.....	12
b.	Recognition Primed Decision Model.....	14
C.	COGNITIVE ARCHITECTURE MODULE	16
1.	Percept Umpire	17
2.	Agent Object	17
3.	Perception, Attention, Working Memory and Situation Formation	17
4.	Meta-Cognition and Long-Term Memory	17
5.	Action Selection.....	18
6.	Communication and Effects of Trust	19
III.	ANALYSIS OF DECISION METHOD AND TRUST EFFECTS	21
A.	DESIGN PARAMETERS.....	21
B.	TEST SCENARIO	23
C.	OUTPUT PROCESSING.....	27
D.	RESULTS – SINGLE AGENT SCENARIO	28
1.	Civil Security Issue Stance	28
2.	Effect of Initial Stance and OAB	29
3.	Effect of Discount Factor and Size of Dataset	30
E.	RESULTS – TWO-AGENT SCENARIO.....	32
1.	Civil Security Issue Stance	32
2.	Decision Method and Action Selection.....	33
3.	Homophily and Communications.....	35
F.	RESULTS – THREE-AGENT SCENARIO	36
1.	Civil Security Issue Stance	36
2.	Decision Method and Action Selection.....	38
3.	Homophily and Communications.....	39
IV.	FURTHER TESTING AND EVALUATION.....	41
A.	DESIGN PARAMETERS.....	41
B.	TEST SCENARIO	42

C.	OUTPUTS	44
D.	RESULTS.....	45
	1. Civil Security Issue Stance	45
	2. Decision Method and Action Selection.....	48
V.	CONCLUSION	53
A.	EFFECTS OF DECISION METHOD	53
B.	EFFECTS OF TRUST	54
C.	OTHER FACTORS.....	54
D.	TRACEABILITY OF ENTITY BEHAVIOR.....	54
E.	FUTURE WORK AND RECOMMENDATIONS	55
	LIST OF REFERENCES.....	57
	INITIAL DISTRIBUTION LIST	61

LIST OF FIGURES

Figure 1.	Theory of Planned Behavior (From Ajzen, 1991).	9
Figure 2.	Cognitive Architecture Components (From Yamauchi, 2012).....	16
Figure 3.	Action Selection Process (From Yamauchi, 2012).	18
Figure 4.	Civil Security Stance over Time - RPD Method.	29
Figure 5.	Time Taken to Reach Steady State Outcome in Issue Stance for Different Discount Factor Settings.	31
Figure 6.	Effect of Discount Factor and Number of Respondents on Civil Security Issue Stance.....	31
Figure 7.	Civil Security Issue Stance for 2-Agent Scenarios.	32
Figure 8.	Experience Level Heatmaps over Time.....	33
Figure 9.	Expected Utility of Infrastructure-related Actions.....	35
Figure 10.	Communications Acceptance/Rejection Rate.....	36
Figure 11.	Civil Security Issue Stance for 3-Agent Scenarios.	36
Figure 12.	Effect of Trust on Deviation in Issue Stance.....	37
Figure 13.	Entity Experience over Time.....	38
Figure 14.	Communications Acceptance/Rejection Rates Between Entities in 3-Agent Scenario.....	39
Figure 15.	Map of Area of Operations (From Yamauchi, 2012).....	42
Figure 16.	Civil Security Issue Stance for Different Initial Stance Levels.....	45
Figure 17.	Civil Security Issue Stance for Initial 50% Adequate.	46
Figure 18.	Distribution of Outcomes - Civil Security Stance at Day 360.....	47
Figure 19.	Infrastructure Node Visitation Outcomes and Effects.	48
Figure 20.	Infrastructure Node Visitation Rates and Outcomes.....	49
Figure 21.	Expected Utility of Infrastructure-related Actions in 6-Agent Scenario.	50

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

Table 1.	Social Dimensions & Categories in Helmand Province Population Narratives (From Hudak & Baez, n.d.).....	11
Table 2.	Input Parameters for six Basic Test Configurations.....	22
Table 3.	Summary of Design Factors and Settings.	26
Table 4.	Description of Key Parameters Measured.....	27
Table 5.	Effect of Trust on Range and Deviation of Issue Stance.	37
Table 6.	Design Points for Final Run.....	41
Table 7.	Definitions for Infrastructure Operation States.....	43
Table 8.	Description of Additional Key Parameters Measured.	44
Table 9.	95% Confidence Interval Levels of Civil Security Stance at Day 360 (Combined Mean across all Entities in Scenario).	47

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

AI	Artificial Intelligence
CF	Coalition Forces
CG	Cultural Geography
COIN	Counter-Insurgency
DoD	Department of Defense
DP	Design Point
EL	Exploration Learning
HA	Humanitarian Assistance
HSCB	Human Social and Cultural Behavior
IW	Irregular Warfare
M&S	Modeling and Simulation
MSCO	Modeling and Simulation Coordination Office
OAB	Observed Attitude & Behavior
OOTW	Operations Other Than War
PEO STRI	Program Executive Office for Simulation, Training & Instrumentation
RPD	Recognition Primed Decision
SSTR	Security, Stability, Transition and Reconstruction
TpB	Theory of Planned Behavior
TRAC-MTRY	TRADOC Analysis Center – Monterey
TRAC-WSMR	TRADOC Analysis Center – White Sands Missile Range
TRADOC	Training and Doctrine Command

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGMENTS

This thesis would not have been possible without the many people who provided invaluable advice and assistance to me throughout the study. I would like to express my heartfelt gratitude for their patience, support and guidance.

My Thesis Advisor, Dr. Chris Darken, was the source of inspiration and motivation for my venture into this area of research. His teachings and in-depth knowledge of the field helped me significantly, and his advice over the course of testing and experimentation provided key insights and guidance that helped shaped this work.

The members of TRAC-Monterey provided invaluable support and coaching for me as we worked to understand the programming wizardry inside the Cultural Geography Model. My Second Reader, Lieutenant Colonel Jason Caldwell, provided me with a deeper understanding and sense of purpose of the work, and helped to broaden my perspective of the inner workings and applications of the Model in operational contexts. Mr. Harold Yamauchi, the programming guru and expert on the CG Model, was instrumental in the scenario creation, manipulation of design inputs and generation of data outputs. His hard work and expertise provided the foundation for all the testing and experimentation that we were able to achieve with the code. I am also grateful to Lieutenant Colonel Jonathan Alt, for giving me this valuable privilege and opportunity to work with TRAC-Monterey; and to MAJ Francisco Baez for his advice and knowledge sharing.

Last, but certainly not least, my sincere appreciation also goes to my wife and son for their support and understanding as they endured my long hours at school and late nights at home.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

A. BACKGROUND

In most modern defense-related ecosystems in the world today, Modeling and Simulation (M&S) has established itself as an effective and resource-efficient tool for training and preparation of military operations and other undertakings. The U.S. Department of Defense (DoD) Modeling & Simulation Coordination Office (MSCO) recognizes that “M&S is an enabler of warfighting capabilities. It helps to save lives, to save taxpayer dollars, and to improve operational readiness” (MSCO, 2012). Wargaming is one common application that allows planners and analysts to gain insight on likely combat outcomes, challenges and potential pitfalls, and other unintended consequences that cannot be captured by traditional analysis methods. In such applications, a key success factor is the ability to maintain an extensive database of fully or semi-automated entities that represent actors within the scenario, and these entities need to have the ability to portray the actions and behaviors of real life combatants. In combat-based models and simulations, relatively realistic portrayal of soldiers and units can be attained through reference to doctrine and tactics, which dictate rules for how the entities would move, interact and react to the situation (Pew & Mavor, 1998; U.S. Army PEO STRI, 2012).

However, in recent times, the spectrum of military operations has expanded tremendously, encompassing missions such as Counter-Insurgency (COIN), Security, Stability, Transition, and Reconstruction (SSTR) efforts, and Humanitarian Assistance (HA) missions. The shift away from conventional conflicts and armed, open fighting between states reflects the changing political and security landscape in the world today. With this, military leaders need the ability and tools to appreciate the planning considerations, courses of actions and challenges in such Operations Other Than War (OOTW) and Irregular Warfare (IW) situations (DoD, 2008; Ng, 2012). In these areas, the changes that military actions bring to the economy, society, and political situation in the area of

operations are often the indicators of mission success (Joint Chiefs of Staff, 1995), and thus the ability to have prior understanding and insights on it is a crucial aspect that needs to be addressed.

Simulating the entities that exist in unconventional environments is complex, as the requirements and challenges for modeling non-combatants and non-traditional combatants such as insurgent fighters are very different. For example, the artificial intelligence (AI) driving the actions of a regular soldier agent may be scripted based on rules of engagement and small-unit tactics; however, the response of civilians in a crowd to the military presence would vary significantly, depending on their demographics, personal circumstances, and perception of the immediate and long-term situation around them.

In this respect, there is a well-recognized need to improve the modeling of realistic human social and cultural behavior (HSCB). This would allow greater fidelity and realism in simulations in the realm of non-lethal operations, where the ability to better captures the “softer” effects of military action and to understand the impact on the population and social structure would be an important contributor to success (Alt, Jackson, Hudak & Lieberman, 2009; Pew & Mavor 1998).

The Cultural Geography (CG) Model developed by the U.S. Army Training and Doctrine Command (TRADOC) Analysis Center – Monterey (TRAC-MTRY) seeks to enhance existing DoD efforts to model the responses of populations and social networks to operations conducted by the military in OOTW and IW campaigns (Alt et al., 2009; TRAC-MTRY, 2009). The CG Model is a multi-agent, discrete event simulation implemented in Java that models populations as entities in a geographical area. The agents, or entities, in the model are based on demographic information defining parameters for their beliefs, attitudes towards other entities, and actions taken. The cognitive architecture module in the CG Model forms the foundation for the artificial intelligence of these entities, and is based on well-studied social theory, concepts and models, such as Icek Ajzen’s Theory of Planned Behavior (TpB), Bayesian Belief Networks, and representation

of homophily and its effects on interactions between entities (Alt et al., 2009; Alt, 2010; Perkins, Pearman & Baez, n.d.).

B. PROBLEM STATEMENT

Currently, the Social Impact Module (SIM) Transition being undertaken by TRAC-MTRY and TRADOC Analysis Center – White Sands Missile Range (TRAC-WSMR) seeks to fine-tune the CG Model to increase its acceptability by the end-users (TRAC-WSMR). One of the possible areas of improvement is to simplify the artificial intelligence and agent behavior in the CG Model so that it is better understood during implementation and use.

The complexity of multi-agent systems like the CG Model, which has many linkages and interactions, makes it realistic as a representation of HSCB, but also increases the difficulty in tracing and understanding the behavior of agents in it, and thus the outcome of the simulation. This thesis seeks to investigate two key aspects in the cognitive architecture of the CG Model. First, the current decision-making process of the entities, which is based on two well-known models – Recognition Primed Decision making (RPD) and Reinforcement Learning (Baez et al. 2010; Ozcan, Alt & Darken, 2011); and second, the trust module within the CG Model, which provides an additional layer of realism (and with it, complexity) by simulating the effect of trust, or the lack of it, between entities in the scenario (Baez et al. 2010; Pollock, 2011).

These components in the cognitive architecture enhance the fidelity of the agent representation as the entities respond based on a greater range of possible options under the effects of the rules that they bring to the model. Individual studies have demonstrated statistically significant contributions of these components to the CG Model (Ozcan et al., 2011; Papadopoulos, 2010; Pollock, 2011). However, in terms of creating a believable, realistic entity that performs on par with end-user expectations, it is worthwhile to consider if similar entity behavior is attained by implementation of a simplified artificial intelligence, i.e., without contributions of varying decision-making methods, or the trust

module. Essentially, an acceptable degree of realism in agent behavior needs to be incorporated in the model, while avoiding an overly prescriptive and cumbersome AI.

C. OBJECTIVES

This study thus aims to isolate and investigate the effects of the decision-making module and the trust module on the outcomes of agent behavior in several test scenarios. As part of the process, it would generate greater insight in tracing the actions of entities, and provide reasonable understanding of the behavior to improve the believability of the model. It would also identify possible areas for simplification in the cognitive architecture, to reduce complexity of the artificial intelligence in the model without compromising on realism.

This thesis seeks to address the following key questions:

1. What significant effects do the decision making and trust components provide in the existing cognitive architecture, and do these perform as expected / desired?
2. Can a simplification of the cognitive architecture provide a reasonable behavior for agents in the CG Model that is comparable with that of the existing framework?

It is envisioned that the experimental design, scenario development and data generated from the study will provide ample references for a better understanding of agent behavior in the CG Model. The study will thus facilitate fine-tuning of the CG Model (in particular the cognitive architecture) towards meeting the requirements of the end-users for the CG Model, as part of the Social Impact Module Transition.

D. METHODOLOGY

The initial thrust of this study was to isolate the components in the cognitive architecture that are of interest, and analyze their effects on outcomes and agent behaviors in a simple scenario with one, two or three entities. Only a

small subset of the full capabilities of the CG Model were used, so as not to introduce excessive effects of external factors which were not being tested. In particular, the agent(s) were placed in a specific geographical location, together with an infrastructure node from which they periodically obtain consumable resources. Scripted actions were injected regularly to trigger responses and changes to entity behavior.

The single entity scenario serves to provide insight on the direct relation between the decision-making method and the entity's behavior and eventual outcome of the scenario. The two-entity scenario added the effect of trust, which would be visible in the form of communications between the two agents. The three-entity scenario furthered the analysis with the addition of another agent based on a distinctly different prototype than the original two. This third entity has a lesser degree of homophily to the other two, and thus the effects of trust and interactions with other agents or the environment would be dissimilar.

This initial analysis measured outcomes in terms of change in population stance, frequency of communications between entities, choice of decision-making method, and the effects of action selections on agent attitudes and stance. Overall, it provided insight on the direct effect that the decision methods and trust have on agent behavior and scenario outcome.

The results of the initial analysis provided the basis for the scenario development of the subsequent set of experiments. The scenario complexity was increased to create a more realistic depiction of a plausible, real-world situation. Six agents and 2 infrastructure nodes were placed in separate geographical locations, but within range of communicating with and reaching each other. Several revisions to the scenario parameters were tested in order to identify one that would best exploit and bring out the differences in the various configurations of the cognitive architecture. The final set up was one in which the infrastructure nodes were initially insufficient to supply the requirements of the agents, but a scripted action was introduced to occur after some time, to improve the state of infrastructure. The intent was to trigger changes in agent behavior after the

occurrence of the scripted action, and identify the variations in response for agents reacting based on the different decision methods and effects of trust.

The data from the initial experimental runs and the various revisions leading up to the final run was analyzed to generate a statistical comparison of the outcomes from the basic decision making methods, with and without trust, compared to the existing cognitive architecture framework in which entities can choose between RPD and Reinforcement Learning, under the effects of trust.

II. OVERVIEW OF THE CULTURAL GEOGRAPHY MODEL

A. DEVELOPMENT

The ‘Representing Urban Cultural Geography’ project was conceptualized in 2006 as an initial prototype for a simulation of a population in a social network (Alt, 2010; Baez et al., 2010; TRAC-MTRY, 2009). Continued work over the next few years saw its development through various forms, with more components and features adding to the depth and complexity of the model, such as inclusion of entity actions (e.g., insurgent activity), representations of resources and infrastructure nodes, communications, and improvements to agent behavior modeling (Alt et al., 2009; Perkins et al., n.d.). The implementation also evolved from its earlier usage of the Pythagoras 2.0 agent based combat model (Ferris, 2008; Seitz, 2008) to its current form, which utilizes the SimKit Discrete Event Simulation in Java (Alt, 2010; Buss, 2011). A key feature of the model is its framework to allowing modules to ‘plug-and-play’ into the program (Alt et al., 2009), allowing flexibility and increased functionality. Two recent CG model developments are of relevance to this thesis—first, the use of a Reinforcement Learning based method for agent action selection (instead of a previous Bayesian network representation) (Yamauchi, 2012); and second, the implementation of a “trust” module that adds onto existing agent behavior. These two components are described in further detail later in this chapter.

As with all models, the intent for the CG Model is not to create a perfectly realistic representation of the world in order predict with absolute certainty what would happen in any given scenario—that would clearly be impossible to achieve. Rather, it provides a framework for analysts and planners to understand a situation and experiment with courses of action and alternatives to assess viability, possible outcomes, and potential pitfalls.

B. UNDERLYING CONCEPTS AND THEORIES

The representation of any real world process or phenomena as a model is intrinsically not an easy task. This is especially true in military and HSCB-based applications where there are a vast number of actors/objects, complex interactions, and lack of well-defined relationships and rules governing causes and effects. In order for the model to perform well, it must produce outputs that are rational and believable with respect to its intended purposes and areas of usage. In the field of HSCB modeling, this can be achieved by building the simulation based on theories in social science and psychology, along with clear understanding of the structure of organizations and demographics of populations being represented (Pew & Mavor, 1998). The CG Model is an example of this, as it is based on well-studied concepts and theories creating a rational and understandable framework for the representation and study of military operations in IW. A brief look at some of the underlying concepts and theories used in the CG Model follows.

1. Theory of Planned Behavior

Icek Ajzen's Theory of Planned Behavior serves as the basis for a core component in the CG Model. This theory attributes a person's intentions and behaviors to three key factors: his attitude towards the behavior, the subjective norms associated with that behavior, and his perceived behavioral control (Ajzen, 1985; Ajzen, 1991). *Attitude towards the behavior* describes the individual's own assessment of the behavior, for example if a person is in favor of always returning to the same provider to obtain a particular resource or commodity. The *subjective norm* brings out the social dimension as it represents the degree to which there is external influence (such as from peers and the community) towards the behavior, for example if a person's local community utilizes a particular other resource provider and pressures him to do the same. The *perceived behavioral control* gives a measure of how easily the individual believes he can carry out the particular behavior, for example if he has the ability

to make the switch to a new resource provider. Ajzen postulates that the combination of these three independent factors determines the individual's intention to behave in a particular fashion, and that the intention and perceived behavioral control in turn determine the actual behavior adopted (Figure 1).

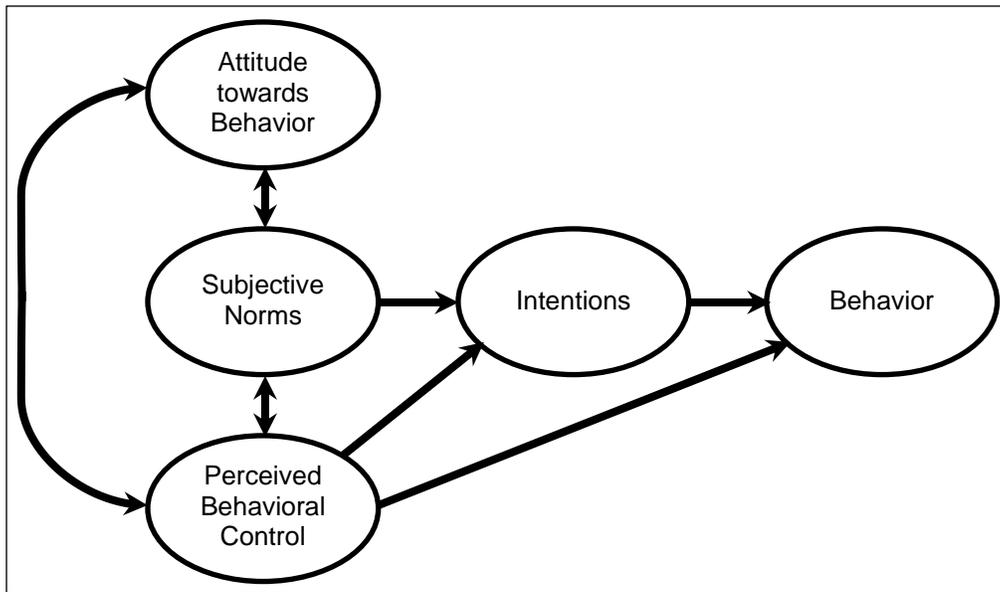


Figure 1. Theory of Planned Behavior (From Ajzen, 1991).

Within the CG Model, these three factors apply to each entity in any given scenario, and are quantified to derive a value for each behavior that the agent may choose. The *attitude towards behavior* is influenced by the agent's demographic stereotype and perception of issues relating to that behavior, the *subjective norm* is determined from the behavior of neighboring agents, and the *perceived behavioral control* is determined from the degree that a selected behavior brings about the agent's desired effect (essentially, a measure of success of behavior choices). User-defined weights are applied to the calculated values of the three factors, and the weighted sum is then used the measure of reward gained from a particular behavior (Yamauchi, 2012), as shown in the formula:

$$v_r = w_A v_A + w_N v_N + w_C v_C$$

where

v_r = reward value of behavior

w_A = weight of Attitude towards Behavior

w_N = weight of Subjective Norms

w_C = weight of Perceived Behavioral Control

v_A = value of Attitude towards Behavior

v_N = value of Subjective Norms

v_C = value of Perceived Behavioral Control

2. Narrative Paradigm

The Narrative Paradigm (Fisher, 1984) provides the logic through which populations in a real-world area of interest are converted to agent representations in the CG Model. Fisher's work proposes that an individual's experiences in life form a collection of narratives that describe his culture and character, shapes his perspective of the world, and affects how he responds to events and interacts with others around him. As such, the narrative account can be used as a comprehensive and credible data set for the purposes of classifying population as different entities, each with its own unique demographic traits and stereotypes for responding to the environment. The CG Model directly implements this by having each entity represent a subset of the population in the area of interest, with the entities ranging from a single individual, to a small group or entire community. Input parameters that are required by the simulation to adjudicate interactions and behavior of agents are then derived from their respective narratives and demographic traits. Table 1 lists the social dimensions and categories for the Afghan Helmand Province data, which was used in this study (Hudak & Baez, n.d.).

Social Dimension	Categories
Family Status	Inherited
	Achieved
	Unemployed
Ethno-Tribal Affiliation	Pro-Government
	Passive
	Marginalized
Disposition	Urban
	Rural
Political Affiliation	Fundamentalist
	Moderate
	Secular
Age	Military Age Male
	SpinGiri ¹

Table 1. Social Dimensions & Categories in Helmand Province Population Narratives (From Hudak & Baez, n.d.)

An entity stereotype is determined by a combination of traits from the list above that forms its demographic profile, along with the initial data of the entity’s attitude and beliefs towards other entities and stance on pertinent issues in the scenario, such as the adequacy of Civil Security in the province.

3. Homophily

The concept of homophily is closely tied to modeling interactions between different population groups in the CG Model. Homophily refers to the similarity between individuals and affects the likelihood that two parties would associate and interact with each other. Its effect is most visible in social network contexts, where similarities and differences in demographic traits and social factors have a pronounced effect on the number and extent of links between people (McPherson, Smith-Lovin & Cook, 2001). This suggests that the effects of

¹ “Spin Giri” is a term referring to senior males who are typically past the traditional warrior/military age, are influential and likely to be local decision makers or hold other positions of tribal leadership (Hudak & Baez, n.d.).

homophily can significantly influence the behaviors of individuals and outcomes of scenarios.

In the CG Model, similarity between entities is determined in accordance with this concept of homophily. The stereotypes (i.e., demographic traits) and geographical proximity of entities are the main factors in the computation, which generates a *homophily link weight* value for each entity pair in the scenario. This link weight is utilized to determine likelihood of communication between the entities, and would affect the sharing of information percepts in the scenario (Alt et al., 2009).

4. Decision Making and Learning

The process of making decisions is a key aspect of human behavior that is modeled in the CG Model. Two main concepts are implemented in the action selection component of the cognitive architecture—the Reinforcement Learning model and the Recognition Primed Decision model.

a. Reinforcement Learning

Reinforcement Learning is a technique of machine learning that determines how agents should act in a situation to generate an optimal overall outcome, based on a specified measure of the estimated value of each possible action. In a given environment, an agent receives information percepts that determine which state it is in, and selects an action from a set of possible options (Russell & Norvig, 2010). The resultant transition to a new state is assessed based on a predefined set of rules, typically in the form of some immediate reward given to the agent. By determining the overall value of each state-action pair (i.e., of choosing a particular action when in a particular state), the agent can make decisions that will allow it to gain the most benefit, or expected utility. The Q-Learning algorithm (Watkins, 1989; Watkins & Dayan, 1992) is implemented in the CG Model. This technique allows the agent to compute and iteratively update the expected utility of actions based solely on the rewards received from them,

and not requiring the environment to be explicitly known, which is well suited for typical scenarios in the CG Model.

Reinforcement Learning provides agents with the ability to adapt well in new situations, where there is a strong impetus for behavior to *explore* possible options and identify the overall optimal course of action. Over time, the value of exploring diminishes as most or all options would have been covered, and the agent can shift its behavior to *exploit* only those actions with high expected utilities. This idea of trade-off exploration and exploitation is well studied; in particular, Ozcan et al. (2011) investigated several techniques for driving agent behavior in the CG model to optimize the balance between them. The action selection process in the CG Model is based on the Softmax method using a Boltzmann distribution, as depicted by the equation:

$$P_i = \frac{e^{E_i/t}}{\sum_j e^{E_j/t}}$$

where

P_i = Probability for selecting action i

E_i = Expected Utility of action i

t = Temperature

The probability of selected a particular action is determined by its expected utility (as compared to that of other actions) as well as a temperature parameter, which influences the exploration-exploitation balance (Baez et al., 2010; Yamauchi, 2012). Thus, an action has a higher probability of being chosen than any other action that has a lower expected utility. In addition, as temperature decreases from its initial value towards zero, the probability of choosing the action with the highest expected utility tends towards one, which gives rise to a purely exploitative behavior.

In the context of the CG Model's cognitive architecture, the Exploration Learning (EL) method² within the action selection module implements this generic reinforcement learning algorithm in accordance with the process developed by Papadopoulos (2010). Papadopoulos identified that the utility-based reinforcement learner was able to function well in the context of selecting the most appropriate action to drive a specified outcome, depending on the settings for parameters such as the initial temperature for the Boltzmann Distribution, learning rate and discount factor of the Q-Learning algorithm and initial expected utilities of actions. These parameters are user-defined values specific to each agent in the scenario, and thus grant the CG Model great flexibility for customization of agent reinforcement learning behavior.

b. Recognition Primed Decision Model

Recognition Primed Decision is a well-known model for naturalistic decision-making propounded by Klein (1989). It describes the theoretical process by which humans are able to make rapid assessment of a situation and come to a good decision without the need for extensive analysis to identify alternatives and then to compare the possible options to deal with the scenario. Klein noted that such behavior could be observed in experienced decision-makers in operational settings, such as firefighter commanders and small unit leaders in the military (Klein, Calderwood & Clinton-Cirocco, 1986; Klein, 1989; Klein, 1999). The RPD model suggests that in complex or time-constrained situations, such experts in their field are able to recognize cues and patterns that allow them to identify an effective course of action quickly, and that this technique would surpass a more deliberate, analytical approach in dealing with the situation.

In the CG Model, the implementation of the RPD model is largely based on the reinforcement learning technique described earlier. During a simulation run, agents will initially utilize the EL method and choose actions in an

² The term "EL" is used here-on to denote the *implementation* of the reinforcement learning algorithm in the CG model. This maintains consistency with the method name used in the CG Model source code and concept diagrams.

almost random manner (assuming that the initial expected utilities of actions are fairly similar). The number of times that the agent has taken any particular action is recorded, and compared to a user-defined minimum threshold, which dictates the number of times that an agent needs to perform each possible action before it is deemed to have sufficient experience. Upon reaching this threshold, the agent will adopt the RPD method of action selection, in which the action with the highest expected utility will always be selected during the decision making process (Yamauchi, 2012).

There are limitations in such an implementation—in particular, it does not capture some characteristics of the RPD model as described by Klein. The implementation in the CG Model is essentially a ‘greedy’ approach of reinforcement learning, where an agent has had the ability to explore various options in the environment before making a decision. In contrast, for a pure RPD approach, this benefit of time and knowledge of action-reward history may not be available to the decision maker. Rather, an agent having made no prior action selections in a particular scenario or environment (and thus having no corresponding estimates of expected utilities of possible actions) would have to decide its course of action based on the limited set of percepts it receives, using other knowledge such as its prior experience and long term memory. In addition, a decision maker in the RPD model would possess the pre-requisite ability to recognize changes in situation and discard previously adopted courses of action that are no longer effective (Klein, 1989; Klein, 1999). The implemented method does not allow agents to have such versatility, thus limiting their ‘expertise’ to situations that are relatively static. Significant changes in a scenario would likely not result in a responsive change of agent behavior once it has adopted RPD, as it would require time for the expected utility of the selected action to drop (until it is no longer the ‘best’ action) before the agent chooses another action.

The RPD model suggests that complex underlying thought processes are involved. For example, picking up cues from a situation (that may only be perceptible to experts but not novices); recognizing patterns that

resemble previously encountered situations; and rapid mental run-through of a possible action to determine its feasibility on its own (as opposed to comparing it against a set of alternatives). These processes cannot be easily incorporated into the existing cognitive architecture of the CG Model, as it could require extensive restructuring of the framework, such as distinguishing between percepts received by expert entities (versus novice entities). This would better represent the significant differences in the performance characteristics of experts in a particular field (Proctor & Zandt, 2008), and thus better suit the implementation of a RPD model. Furthermore, it could require the introduction of larger and more complex long-term memory structures that can be used to compare past scenarios and experiences of an agent against a new situation in which it has limited percepts and situational awareness. Given the constraints in the cognitive architecture framework and the limitations of the current implementation, the RPD method in the CG model is an imperfect but necessary substitute for an actual RPD model.

C. COGNITIVE ARCHITECTURE MODULE

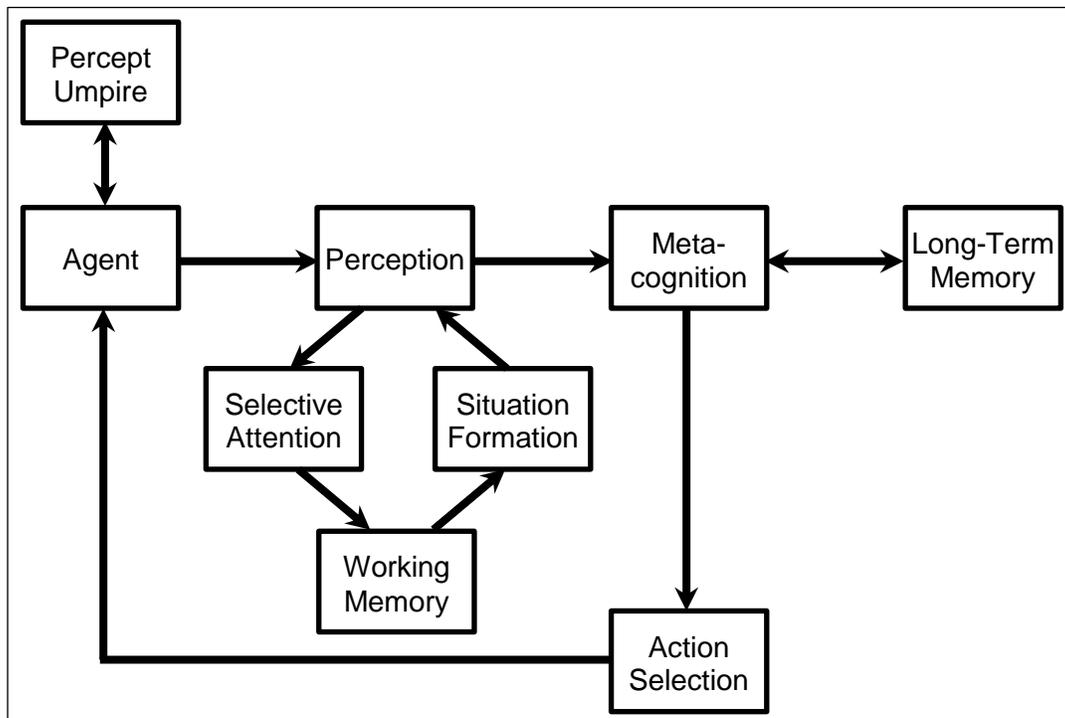


Figure 2. Cognitive Architecture Components (From Yamauchi, 2012).

The main components of the cognitive architecture module are shown in Figure 2, and their functions are described below.

1. Percept Umpire

The Percept Umpire acts as the 'sensor' for agents in the CG model. It receives information from the environment and entities in the model, such as changes to the state of infrastructure nodes, actions carried out by entities and consumption of resources by entities. These are scheduled as percept arrival events for the entities that are supposed to receive them.

2. Agent Object

The Agent component manages the actual state of entities in the CG Model, and is responsible for scheduling events such as performing actions, consuming resources and passing on percepts to the environment and other entities (through the percept umpire).

3. Perception, Attention, Working Memory and Situation Formation

When the entity receives percepts via the percept umpire, the Perception component of its cognitive architecture manages this incoming information, such as monitoring if the agent has the selective attention capacity to accept the information; checking the percept for relevancy and storing it in the working memory of the agent; and using this to schedule the meta-cognition events which are the precursors to the entity's decision making and action selection processes.

4. Meta-Cognition and Long-Term Memory

The meta-cognition and long-term memory components represent the entity's comprehension and assessment of its situation. Key events such as changes in attitude towards other entities or issues are scheduled within these components. The outcome of these stages is to determine possible courses of action for the entity based on the external situation and its internal motivations,

attitudes and beliefs, and schedule the event for the agent to select a decision-making method and then make a decision.

5. Action Selection

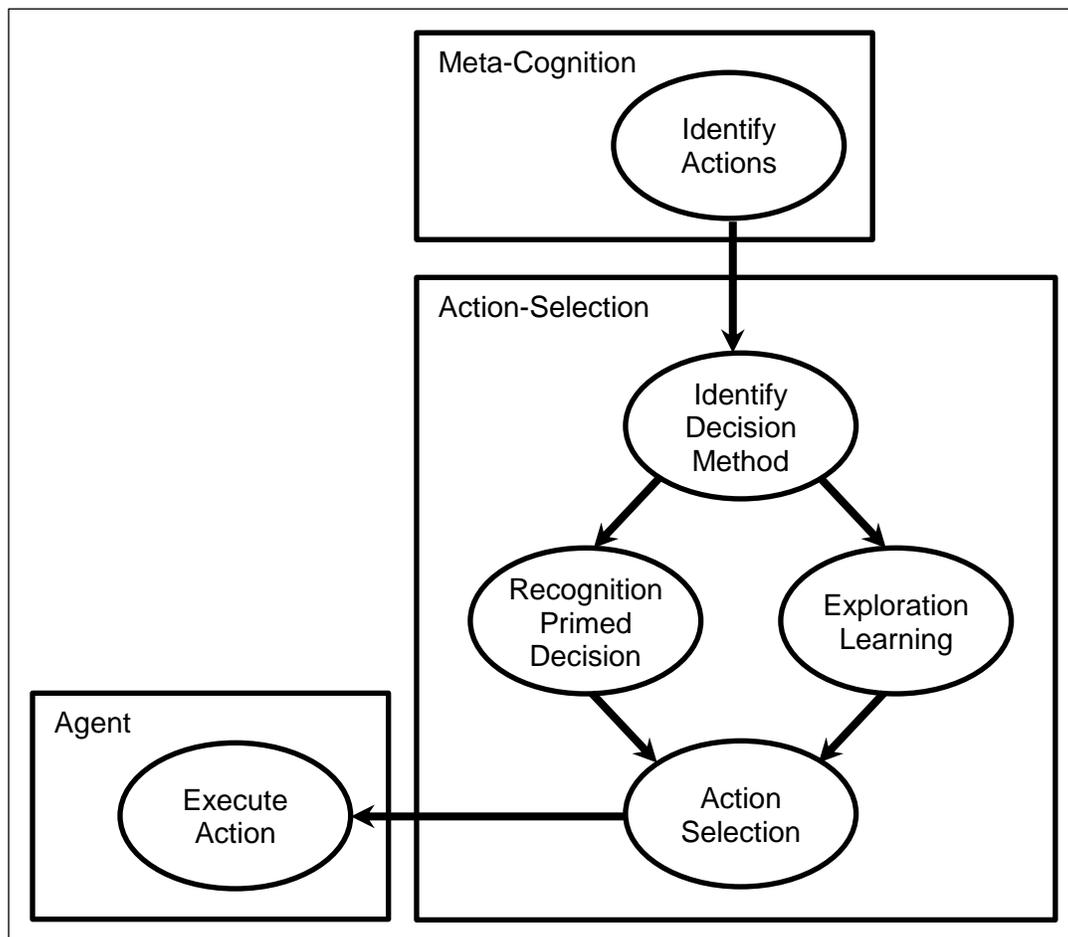


Figure 3. Action Selection Process (From Yamauchi, 2012).

The action selection component (Figure 3) is the main aspect of the cognitive architecture that is studied in this thesis. The process begins with the list of actions received from the meta-cognition component, which determines the type of decision-making method to use—either Exploration Learning (EL) or Recognition Primed Decision (RPD). The event to determine this takes into account the number of times that each possible action has been performed in the past, with the lowest count deemed as the entity's experience. This gives a simple and effective check to assess if the agent has sufficiently sampled all

possible state-action pairs to build an accurate estimate of their expected utilities. Either the RPD method or EL method is scheduled, depending on whether the minimum experience has been reached. Thus, the minimum experience threshold parameter (pre-defined by the user) directly controls the amount of exploration that entities are allowed before they settle in the 'greedy' RPD mode. Once the decision-making method has been determined, the entity selects the appropriate action based on the probabilities evaluated from the range of expected utilities (or, simply selects the action with the highest expected utility in the case of RPD), and schedules the event for it to be carried out.

The action selection process also includes methods to initiate other scheduled events such as scripted behavioral actions and the cancellation of existing actions if necessary. These are methods are not investigated for the purposes of this study.

6. Communication and Effects of Trust

The CG Model simulates the interaction of entities and passing of information as communication actions taken by agents, such as the sending and receipt of percepts between them. This interaction influences the decisions and actions of entities, as it influences the parameters that are passed through their planned behavior process, in particular their *attitudes towards behaviors* and the effect of *subjective norms*. Pollock (2011) developed algorithms for representing trust between entities in a social structure, which aimed to capture additional facets of the relationships and effect of communications between agents.

Scenario designers initialize entities with parameters that determine their frequency of communication with other agents, while their similarity to others (as expressed through the homophily link weights) influences who they choose to communicate with. The trust filter implemented by Pollock interjected a check into the communication process that measures the level of trust between two communicating agents. The parameters for initial trust and changes to trust levels during run-time are defined in the scenario set up. With this trust filter,

entities will still receive, but not accept or process, information received from agents that do not satisfy minimum trust requirements (Yamauchi, 2012). Pollock (2011) noted that inclusion of trust into the interactions reduced the rate at which agent changed their beliefs to align themselves with others. This study will look further at the effect on the overall scenario outcomes, as well as possible influences in conjunction with the choice of decision-making method.

III. ANALYSIS OF DECISION METHOD AND TRUST EFFECTS

A. DESIGN PARAMETERS

The experimental set up was designed to test two main aspects in the cognitive architecture of the CG Model—the decision making method, and the effect of trust. This corresponds to the following six basic test configurations:

1. Recognition Primed Decision only, without the effects of trust.
2. Recognition Primed Decision only, with the effects of trust.
3. Exploration Learning only, without the effects of trust.
4. Exploration Learning only, with the effects of trust.
5. Selection of either Recognition Primed Decision or Exploration Learning, without the effects of trust.
6. Selection of either Recognition Primed Decision or Exploration Learning, with the effects of trust. This is the typical configuration that is used in the current CG Model.

The tests were conducted using the Tactical Wargame 2011 (Revision 1160) version of the CG Model, as well as a modified variant of this version for the RPD only cases, in which the EL method of action selection was disabled. Entities in the RPD only variant would consistently choose the action that has the highest expected utility. This implementation serves to remove or reduce the ability of agents to gradually explore possible options and iteratively evaluate the expected utilities of all actions, and thus mimics human behavior in accordance with Klein's model of RPD. However, it is still limited by the inability to duplicate the process of rapidly assessing a new situation and selecting an effective solution based on one's expertise. The test configurations in which entities only use the Exploration Learning method were created by implementing a very high minimum experience threshold of 1000. This meant that the agents were forced to consistently choose the EL method over RPD, as the scenario run times were

not long enough for them to have attempted all possible actions at least 1000 times each. The baseline configuration where entities could adopt either RPD or EL was set up using a minimum experience threshold of five.

The trust effects were tested by disabling the calculations of trust in code for the relevant configurations. The result of this is to prevent entities from performing checks that would disregard communications from senders whom they did not trust.

All other input parameters that are required for proper functioning of the cognitive architecture (in particular, for the Q-Learning Algorithm, Softmax algorithm, behavior utility calculations and trust module) were kept constant across the 6 test configurations. Table 2 summarizes the key input parameter settings that were used.

Configuration	1	2	3	4	5	6
Decision Method Settings	EL method disabled		Min Experience Threshold = 1000		Min Experience Threshold = 5	
Trust Filter Settings	Off	On	Off	On	Off	On
Reinforcement Learning Parameters	Initial Temperature = 0.1 Discount Factor, Lambda (λ) = 0.01 or 0.1 (see below)					
Behavior Parameters	Weight of Attitude towards Behavior = 0.3 Weight of Subjective Norms = 0.3 Weight of Perceived Behavioral Control = 0.3					
Trust Parameters ³	Default Trust = 0.5 Learning Rate = 0.8 Discount Factor = 0.3 Trust Temperature = 0.5					

Table 2. Input Parameters for six Basic Test Configurations.

³ Pollock (2011) provides a detailed investigation of the effects of these parameters, which are used in the algorithms pertaining to the reinforcement learning of trust, and affect the rate at which entities' trust fluctuate during the scenario runs.

In addition to the six test configurations, three other factors were varied for the initial set of tests: (1) the Reinforcement Learning Discount Factor, Lambda (λ), (2) the effect of scripted actions taking place during the scenario, and (3) the initial belief and issue stance of entities in the scenario. These factors had earlier been studied as part of the ongoing testing and evaluation by TRAC-MTRY, and were incorporated in the initial run to extend the number of data points over which the basic configurations could be tested.

The reinforcement learning discount factor (λ) was tested at two levels (0.01 and 0.1). The former corresponds to behavior that favors short term rewards, as the value of rewards (i.e., their contribution to expected utility of an action) diminishes more rapidly with time, while the latter corresponds to behavior that favors longer term rewards.

The effect of scripted actions was set to be either positive or negative, while the initial belief and issue stance of entities was varied over 14 possible cases. Further elaboration of these two factors is provided in the next section.

B. TEST SCENARIO

For the purposes of the initial run, a simplistic test scenario was used in order to minimize interactions from other components in the CG Model, and allow the effects of the test configurations to be isolated. This test scenario was developed based on the Helmand Province Case Study developed by the IW Study Team at TRAC-MTRY (Baez et al., 2010; Hudak & Baez, n.d.). The study encompassed several districts in the province, and generated a significant amount of data and analysis pertaining to the population demographics and their views three key issues—security, infrastructure and governance. It serves as a well-documented starting point for the purpose of scenario creation in the CG Model by providing rich datasets that facilitate the development and selection of initial parameters, and has been used in several other studies conducted by TRAC-MTRY (Alt et al., 2009; Perkins et al., n.d.; Wiedemann, 2010).

In the test scenario, two identical infrastructure nodes were sited within the area of operation, and constantly provide a consumable resource (electricity) to either one, two or three agents in the scenario. These agents consume the resource at a constant rate, and may carry out the action of visiting the infrastructure nodes to restock their supply as dictated by their behavior.

In the 1-agent and 2-agent cases, the entity prototype was assigned the social dimensions of *Inherited* family status, *Pro-Government* ethno-tribal affiliation, *Urban* disposition, *Secular* political affiliation, and *Spin Giri* age group. This is a typical entity used in the CG Model, abbreviated as I_P_U_S_Sp. In the 3-agent cases, the third entity was assigned social dimensions that were dissimilar from I_P_U_S_Sp – Unemployed, Passive, Rural, and Moderate, and Military age (Un_Pa_R_M_Ma). This distinction reduces the degree of homophily between the third agent and the other entities, to lower their homophily link weights and bring out any differences in behavior due to the effects of trust.

The population stance on the issue of civil security was used as the primary measure of scenario outcome, and the overall effects of the test parameters. This issue stance represents the percentage of the population (more precisely, of the groups represented by each entity in the scenario) who perceive that the level of civil security in the province is adequate. This issue stance is affected by many factors in the model, such as the beliefs of a particular demographic group as determined by their population narrative (e.g., the belief that Coalition Forces are not trustworthy or that the area is not a safe). Also, the occurrence of events during run-time (such as Insurgent or CF activity) and information passed on from other entities during the scenario (Yamauchi, 2012) are significant influences on the issue stance..

In addition, each entity possesses a set of attitudes and behaviors towards certain groups or issues. This is quantified as an *observed attitude and behavior* (OAB), which translates to one of five levels—positive-active (PA), positive-passive (PP), neutral (N), negative-passive (NP), and negative-active (NA). The OAB of interest to this study is that pertaining to the entities' perception of CF

(*OABtowardsCF*). An entity that is positively inclined towards CF but does not actively carry out actions in support of them would have an *OABtowardsCF* value that falls in the range corresponding to positive-passive; an entity that is negatively inclined and is likely to choose actions such as aiding insurgents would have an *OABtowardsCF* in the level of negative-active (Yamauchi, 2012).

Seven different settings were used for the initial belief and issue stance (“casefiles”) of the entities in the test scenario. These correspond a combination of high/low extremes and mid-point levels for these two parameter (issue stance on civil security and *OABtowardsCF*), and are shown in the summary of design factors/levels in Table 3.

In addition, a periodic scripted action was implemented in the scenario, representing the operation of Coalition Forces (CF) within the area that is visible to the agent(s). This scripted action was programmed to have a positive effect on the population stance on the issue of civil security in the area for half of the test cases, and a negative effect for the rest.

A final parameter that was varied was the size of dataset used as input parameters. This represents the sample size of the data collection process that is used to generate the entity stereotypes based on the population narratives. A setting of either 1000 or 100 respondents was used, to verify that reduction of the sample size would not have an impact on the consistency of results or overall outcome of scenario.

With 6 basic configurations – three settings for decision method (RPD / EL / Both) times two settings for trust (ON / OFF) – two settings for discount factor, seven settings for initial belief and stance, two settings for scripted action effect, and two settings for data sample size, a total of 336 design points were generated for the 2- and 3-agent scenarios. One hundred sixty-eight design points were generated for the 1-agent scenarios (as the trust-ON setting is irrelevant in this context). This created a total of 840 design points for the initial run. Table 3 provides a summary of the factors and settings used.

Factor	Number of Settings	Settings
Number of Agents	3	1-Agent: I_P_U_S_Sp_1
		2-Agent: I_P_U_S_Sp_1, I_P_U_S_Sp_2
		3-Agent: I_P_U_S_Sp_1, I_P_U_S_Sp_2, Un_Pa_R_M_Ma_1
Decision Method	3	RPD Only
		EL Only
		Both
Trust	2	On (Not applicable in 1-Agent case)
		Off
Discount Factor	2	0.1
		0.01
Scripted Action Effect	2	Positive
		Negative
Dataset Sample Size	2	100 Respondents
		1000 Respondents
Initial Casefile	7	Civil Security Stance: 100% Adequate OAB towards CF: 99% PA, 1% NA
		Civil Security Stance: 99% Adequate OAB towards CF: 99% PA, 1% NA
		Civil Security Stance: 50% Adequate OAB towards CF: 99% PA, 1% NA
		Civil Security Stance: 50% Adequate OAB towards CF: 50% PA, 50% NA
		Civil Security Stance: 50% Adequate OAB towards CF: 1% PA, 99% NA
		Civil Security Stance: 1% Adequate OAB towards CF: 1% PA, 99% NA
		Civil Security Stance: 0% Adequate OAB towards CF: 0% PA, 100% NA

Table 3. Summary of Design Factors and Settings.

Each design point was replicated 30 times, using a fixed set of 30 random seeds for all design points. The scenario was allowed to run for 140 days (simulation time), to allow sufficient time for trends in the performance measure to be seen, and steady state outcome to be observed.

C. OUTPUT PROCESSING

Dataloggers in the CG Model were used to record pertinent data from the scenario replications during run-time. The key parameters that were measured are shown in Table 4.

Parameter	Datalogger(s) Used	Description
Civil Security Issue Stance	PositionChange-PeriodicDataLogger PositionChange-DataLogger	Each entity's stance on the issue of civil security was recorded on a daily basis to monitor its change over time. Specific events (e.g. receipt of communications) resulting in changes in stance were also recorded.
Choice of Decision Method and Actions	DecisionMethod-DataLogger SelectAction-DataLogger	Every occurrence of the event where an entity chooses a particular decision method (RPD or EL) was logged, along with the entity's level of experience at that time. The action selected as a result of the decision method used, and the expected utility of the action, were also recorded.
Communications	CommCount-DataLogger Communication-DataLogger	All communication events between entities were recorded to keep count of the total number received by each entity, and the number that the entity rejected (due to the trust effects) The trust level between the two entities involved in each communication event was also logged.
Degree of Homophily between Entities	HomophilyNetwork-DataLogger	The homophily link weights between any 2 entities in the scenario were recorded periodically (every 30 days).
OAB	PositionChange-DataLogger	The OAB of entities towards CF was recorded for each event that triggered any changes in the level. This log measured the percentage of the population represented by each entity that fall into each of the 5 levels of OAB. This parameter was tracked for the purpose of cross-referencing with the issue stance, but not used directly as a measure of scenario outcome. ⁴

Table 4. Description of Key Parameters Measured.

⁴ Prior testing and evaluation by TRAC-MTRY had suggested that issue stances were more appropriate and better understood as measures of changes and outcomes in scenarios, compared to OABs. (J. Caldwell & H. Yamauchi, personal communication, July 2012).

Due to the large volume of data generated⁵, a combination of manual and batch-file processing methods were used to organize the outputs into similar dataset groupings. These were further processed with SAS Institute's JMP Pro (version 10) statistical software to consolidate datapoints into relevant parameters, such as mean and variance across replications, trends over time periods in the scenario, and differences between entities and initial casefiles. JMP was also used for the analysis of the data and generation of plots.

D. RESULTS – SINGLE AGENT SCENARIO

The single agent scenario demonstrated the effects of the design factors at the most primitive level. The effects of trust, homophily and communication were not seen in this scenario as there were no inter-agent interactions taking place.

1. Civil Security Issue Stance

Figure 4 shows the trend of civil security stance of the single entity I_P_U_S_Sp in the case where RPD is fixed as the only option for decision making method. The 28 plots depict the differences across the 14 different casefiles (7 variants of initial stance and OAB with 2 settings for the effect of scripted actions) and settings for the discount factor. From left to right, the columns correspond to the casefiles with initial stance of 100% inadequate, 99% inadequate, 99% adequate, 50% adequate with 99% PA, 50% adequate with 50% PA, 50% adequate with 99% NA, and 100% adequate. The upper 14 plots are for the cases where the scripted action has a negative effect on the entity, while the lower 14 are for the cases with a positive scripted action effect. The plots on the first and third rows correspond to the discount factor of 0.01, while the second and fourth rows show trends with discount factor set to 0.1. The change in scenario outcome as a result of the scripted action conforms to

⁵ Eight output files in comma-delimited value format were generated for each design point, corresponding to 6720 data files in total. Each file contained approximately 4200 to 12600 datapoints, depending on the type and frequency of parameters logged.

expected behavior—the shift in entity perception of civil security issue stance is in the same direction as the effect caused by the periodic scripted action for all test cases.

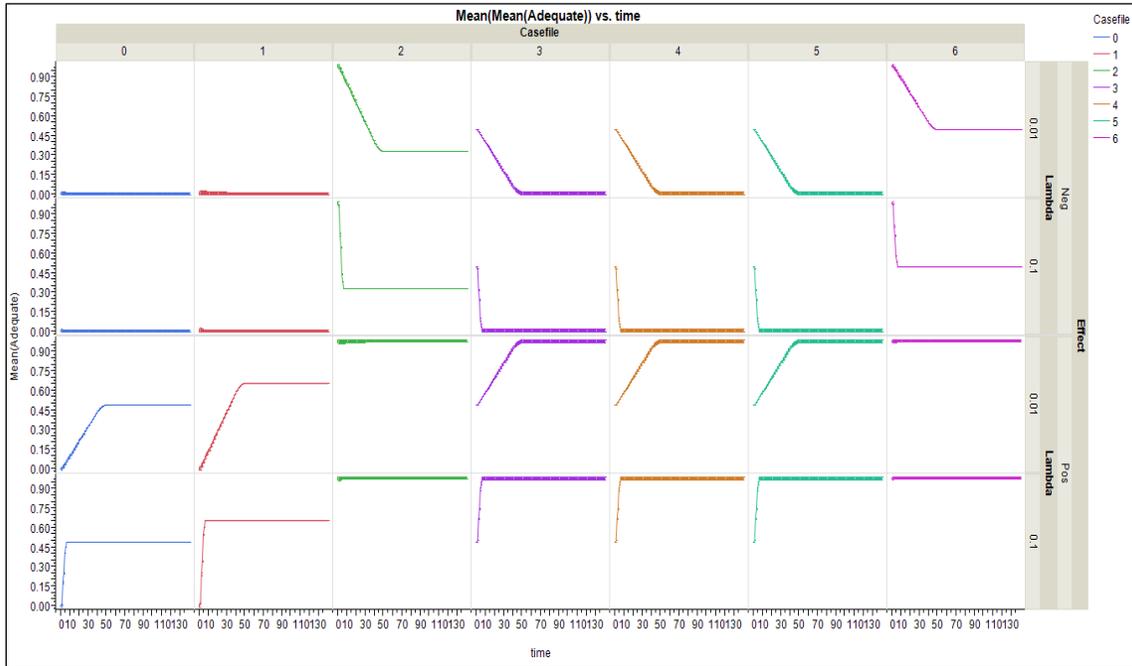


Figure 4. Civil Security Stance over Time - RPD Method.

The variation of both the trend and final state of civil security stance was observed to be unaffected by the decision method adopted by the entity in these test cases. The plots for the settings of EL and BOTH for the decision method were identical to that of the RPD case. This was a clear indication that the decision method was having little or no effect on the final scenario outcome in this set of single agent test cases, which was to be expected, in view of the limited impact that the agent’s action selection had in the simple scenario set up.

2. Effect of Initial Stance and OAB

The initial casefiles used for the entity had a significant impact on the scenario outcome. Comparing the cases of 100% inadequate and 99% inadequate, the difference of just 1% resulted in a significant impact on the final

level of the issue stance, seen in the bottom left most plots of Figure 4. The same effect was noted in the opposite case, where the initial stance was either 100% adequate or 99% adequate. However, from the 3 casefiles where the population started at 50% level of perceived civil security adequacy, it was noted that the initial OAB towards CF did not cause any change in the final outcome of the scenario. These observations point to the importance of the initial data development process in the CG Model, which constructs casefiles and agent prototypes used in any scenario. The effect of initial stance is further studied in the subsequent test scenarios.

3. Effect of Discount Factor and Size of Dataset

A highly notable observation from the single agent dataset was the significant effect of the discount factor setting on the rate of change of issue stance. Comparing across all test cases with a reinforcement learning discount factor of 0.01, the simulation time required for the issue stance to reach its final steady state was between 3 to 6 days. However, with the discount factor set at 0.1, the time taken ranged from 36 to 49 days. Figure 5 shows the distribution of time taken to reach steady state for replications of the test cases based on an initial stance of 50% adequate, with 50% of the population being positive-active towards CF. The final value of the issue stance was unaffected by the different settings of discount factor. However, it was noted that the issue stance at steady state for the case was affected by the size of dataset used (i.e., the number of respondents on which the casefiles were based). Figure 6 shows the combined effect of the discount factor and number of respondents across the 30 replications of the design point.

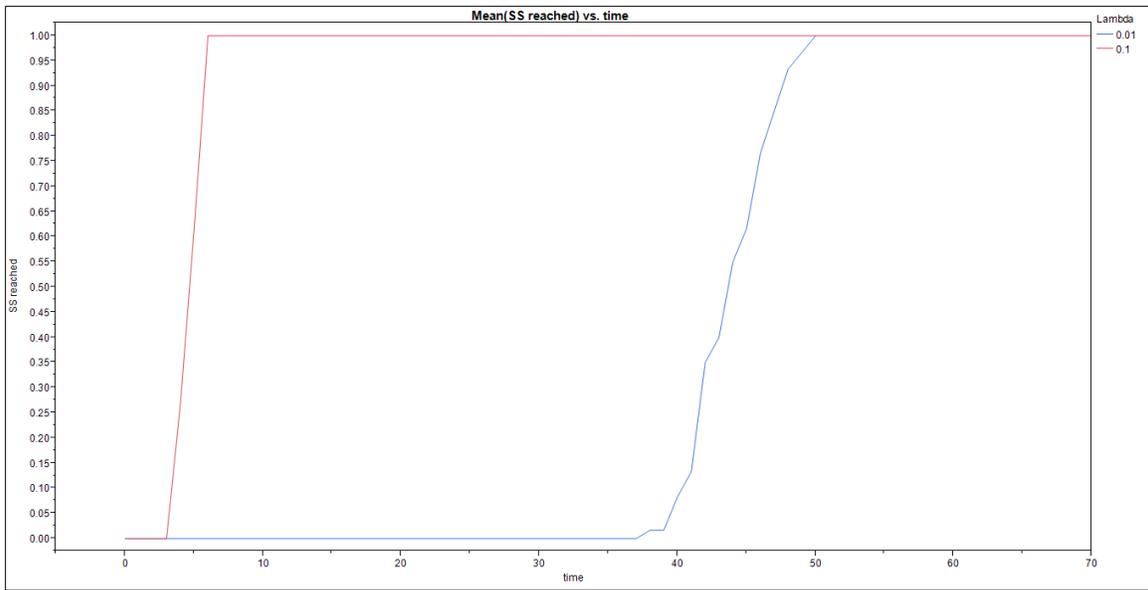


Figure 5. Time Taken to Reach Steady State Outcome in Issue Stance for Different Discount Factor Settings.

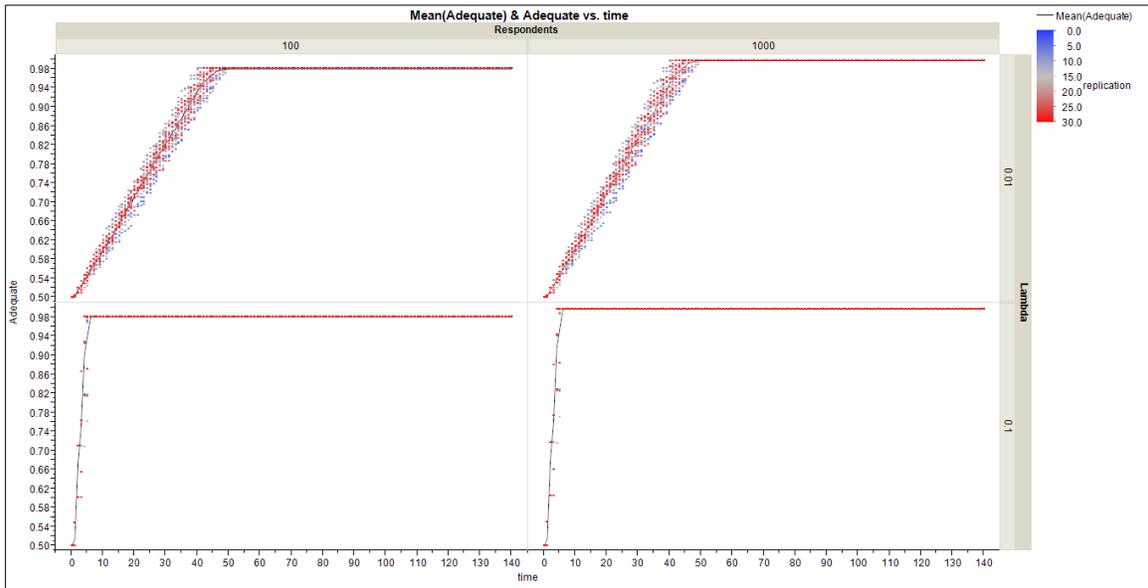


Figure 6. Effect of Discount Factor and Number of Respondents on Civil Security Issue Stance.

E. RESULTS – TWO-AGENT SCENARIO

The results of the two-agent scenario were generally in line with the key observations made from the single agent cases. The data analysis and post processing focused on the design points with the settings of 100 respondents and discount factor of 0.01. This was in consideration of the fact that the cases for 1000 respondents was largely similar to those for 100 respondents, and that the discount factor of 0.1 resulted in behavior (and corresponding scenario outcomes) that shifted too rapidly.

1. Civil Security Issue Stance

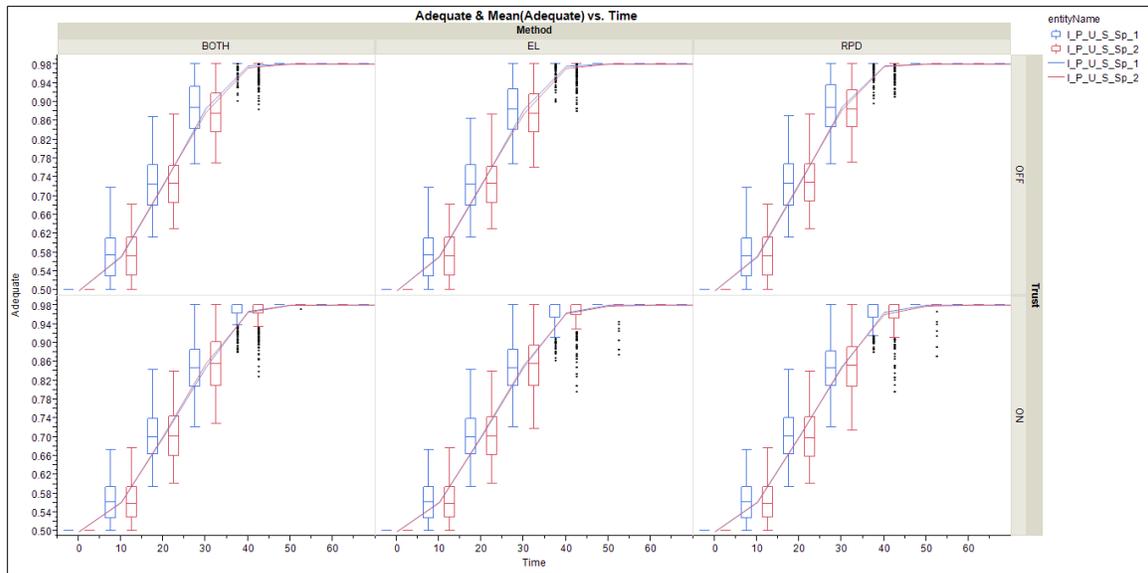


Figure 7. Civil Security Issue Stance for 2-Agent Scenarios.

Figure 7 shows the trend of civil security issue stance over time, for the cases with initial stance at 50% adequacy and positive effect of scripted actions. The stance of both entities remained fairly close to each other throughout the scenario run time, with variations in mean of less than 2% at any point in time. Significant spread was noted across the replications in all six test configurations for the interval in which the stances were shifting from their initial to final states, with a range of up to 22% within each discretized time block of 10 days. The final

outcomes and time to reach steady state were comparable to the earlier single agent test cases, with little variation observed between the different decision methods and effects of trust.

2. Decision Method and Action Selection

The effects of decision-making were studied in detail in the two agent scenarios. Figure 8 is a representative plot of the outcomes of decision-making processes for the 50% initial stance cases, showing the experience levels of the entities over time, across the 30 replications of each design point.⁶

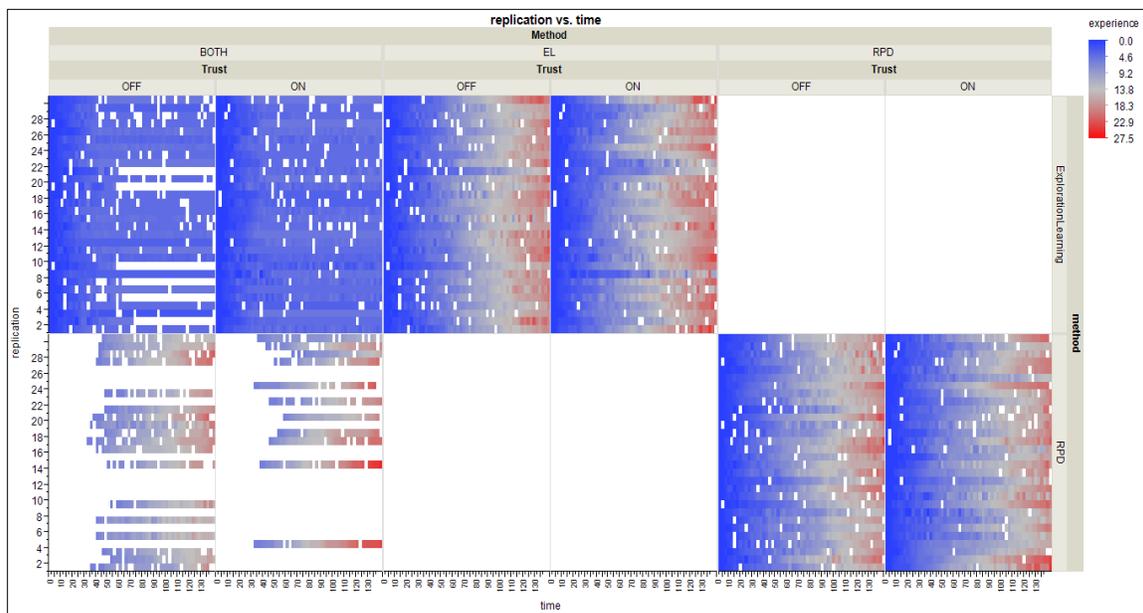


Figure 8. Experience Level Heatmaps over Time

In the design points where the entities could adopt either RPD or EL (heatmaps on left), EL was observed to be the initial choice for decision-making method, as expected. Entity behavior switched to the RPD method for 18 out of 30 replications in the design point where trust was OFF, and 11 out of 30 in the design point where trust was ON. In the cases where EL was maintained throughout the entire duration of the replication, it was observed that the

⁶ Blanks within the plots indicate points in time where the event of selecting a particular decision-making method did not occur, and thus no experience level was logged.

experience level of the entities in those runs remained fairly low throughout the scenario. In contrast, with the design points that only allowed EL (plots in center), entity experience continued to rise to significantly higher levels for the majority of replications. Furthermore, the experience that entities attained was comparable to the cases of RPD method only (plots on right).

The observed trend in experience levels of entities using the different decision-making methods highlights a peculiarity of the current implementation of the cognitive architecture. As the RPD method here is essentially a reinforcement learning based technique with a greedy approach, entities that switch to RPD would always select the action that yields the best return. This would suggest that a certain set of actions would consistently not be chosen, if they were associated with the lowest expected utilities, and thus the experience of entities should remain at that value (of the minimum number of times which those actions had been performed). This is clearly not the case in the data observed, as the RPD only cases showed continued rise in experience level, suggesting that other factors are influencing change in behavior or utility of the actions that would otherwise not be used. The EL behavior seen in the plots appear to conform to expectations, with a gradual increase in experience over time, as the entities would be likely to attempt all actions and thus increase the minimum number of times which each has been chosen. These results suggested the need for further study of the decision method selection process and action selection process.

Figure 9 shows the mean expected utilities of the three possible actions pertaining to infrastructure consumption. Agents are able to choose between using their existing service provider (“Use_Current_Provide”), switching to another (“Seek_New”), or decide not to attempt to restock their resources (“Do_Nothing”). The expected utilities for the actions of seeking a new provider or remaining with their existing ones are expected to be similar in this case, as the nodes available to the entities are essentially identical. The trend of expected utilities over time indicate that entity behavior is reasonable in this case—over time, they would continually make the choice of seek out either infrastructure

node to resupply themselves, instead of doing nothing. However, it is noteworthy that there is no marked difference for the different decision-making methods or trust settings.

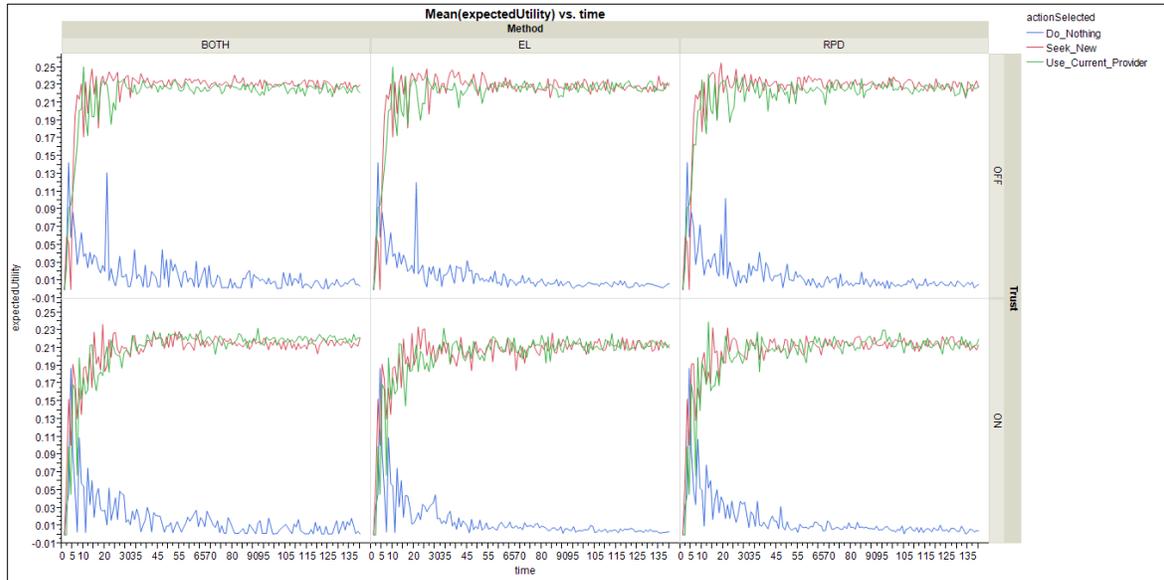


Figure 9. Expected Utility of Infrastructure-related Actions.

3. Homophily and Communications

The homophily link weight between the two entities did not vary with the different decision methods and trust settings. However, the effect of the trust was observed from its effect on communications between the entities. The initial trust level between the entities in these cases was set at 0.5, which rapidly increased to close to the maximum of 1.0 as expected, given the high degree of homophily between them (since they are built on the same prototype). The percentage of communications between the entities that were accepted thus increased over time, from an initial 66% to 87% by the end of the simulation (Figure 10).

	previousTrust	commDecision	
		RECEIVE_ACCEPT	RECEIVE_DONT_ACCEPT
Time	Mean	Row %	Row %
10	0.67295154	66.22%	33.78%
20	0.8845323	81.87%	18.13%
30	0.95944339	87.28%	12.72%
40	0.97715337	87.56%	12.44%
50	0.984732	87.01%	12.99%
60	0.9887895	86.99%	13.01%
70	0.9887569	86.53%	13.47%
80	0.99190848	88.01%	11.99%
90	0.99192274	86.95%	13.05%
100	0.99084663	87.41%	12.59%
110	0.99201641	88.56%	11.44%
120	0.99276436	87.62%	12.38%
130	0.99356156	87.88%	12.12%
140	0.99254373	87.19%	12.81%

Figure 10. Communications Acceptance/Rejection Rate.

F. RESULTS – THREE-AGENT SCENARIO

1. Civil Security Issue Stance

The civil security stance in the 3-agent scenario showed a similar trend over time as that of the 2-agent case (Figure 11).

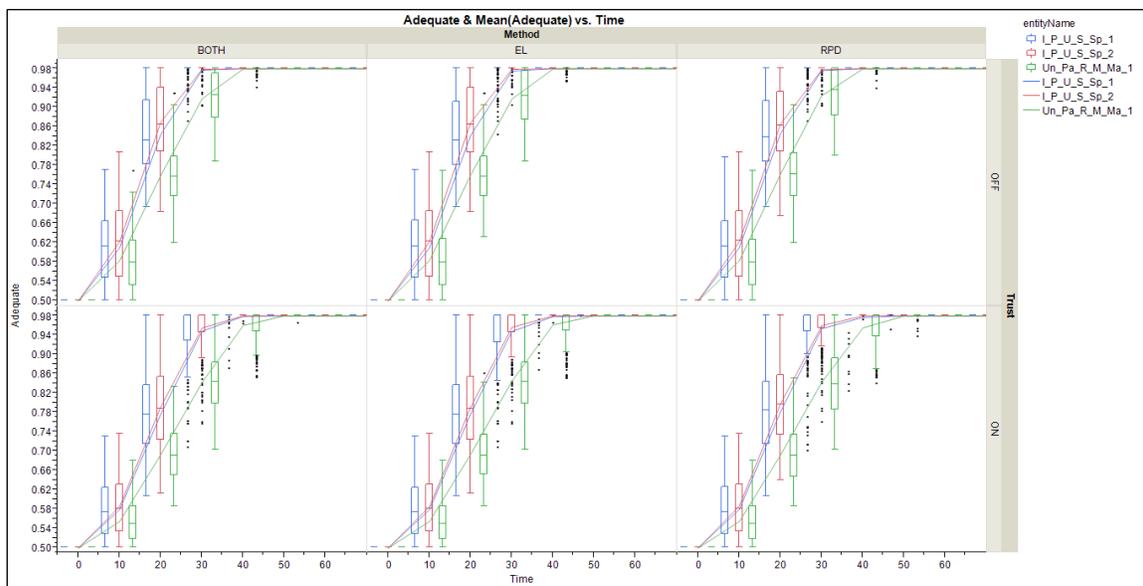


Figure 11. Civil Security Issue Stance for 3-Agent Scenarios.

The new agent, Un_Pa_R_M_Ma_1 demonstrated behavior similar to the original two, but took a longer time to reach its final state in issue stance. The effect of communication was clearly the cause of this behavior—at the 40 day mark, the Un_Pa_R_M_Ma_1 entities in the test cases where the trust module was deactivated had all reached steady state of 98% adequate. In contrast, for the cases with trust on, the mean issue stance in the same time period was 96%, with a 3% standard deviation and range from 87% to 98%. Figure 12 and Table 5 compare the standard deviation of issue stance over time under the effects of trust. The variance is significantly increased for all cases where the trust module is active, but not affected by the decision method used.

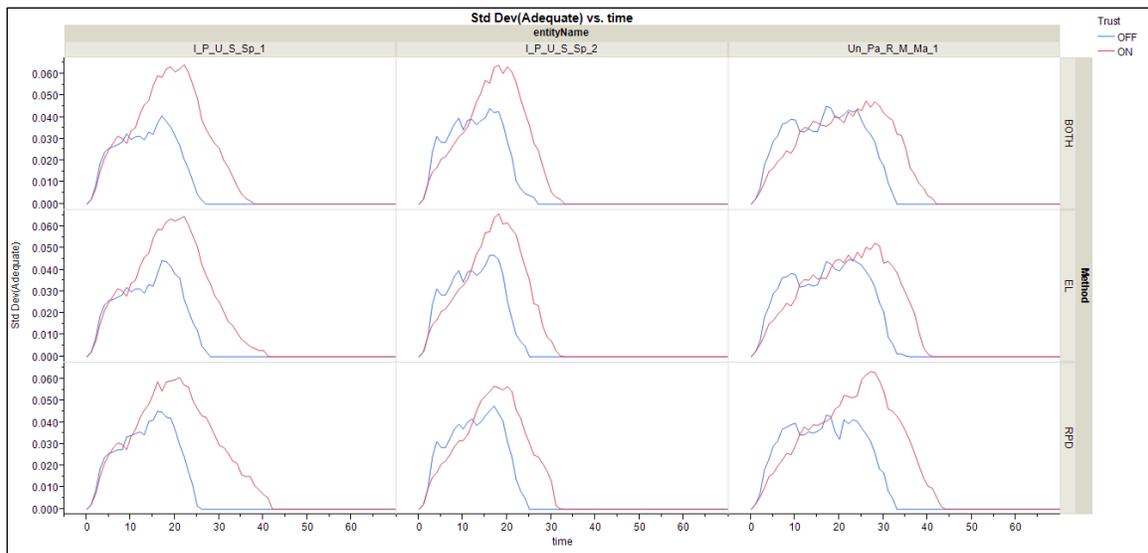


Figure 12. Effect of Trust on Deviation in Issue Stance.

Entity	Trust	Max. Range	Peak Std Dev.	Max. Time to Steady State
I_P_U_S_Sp_1	ON	30.4% (Day 19)	6.5% (Day 22)	Day 43
	OFF	18.4% (Day 15)	4.5% (Day 16)	Day 28
I_P_U_S_Sp_2	ON	27.2% (Day 17)	6.6% (Day 18)	Day 32
	OFF	20.8% (Day 15)	4.8% (Day 17)	Day 27
Un_Pa_R_M_Ma_1	ON	21.5% (Day 26)	6.4% (Day 27)	Day 44
	OFF	18.9% (Day 10)	4.5% (Day 17)	Day 34

Table 5. Effect of Trust on Range and Deviation of Issue Stance.

2. Decision Method and Action Selection

The experience levels of the three entities were comparable throughout the progress of the scenario, and the results showed behavior similar to the 2-agent cases. Additionally, as seen in Figure 13, the trend of experience gain by entities in RPD or EL only modes was distinctly different from the cases where both decision methods were admissible. As before, the expected behavior in EL mode matched the experience trend observed, but that of RPD mode did not. These findings reinforce the notion that the implementation of RPD in the CG Model is in essence a reinforcement learning type approach, but also point out that the process of choosing between EL and RPD alters the behavior of the entities such that the outcome differs from a pure EL or pure RPD scenario.

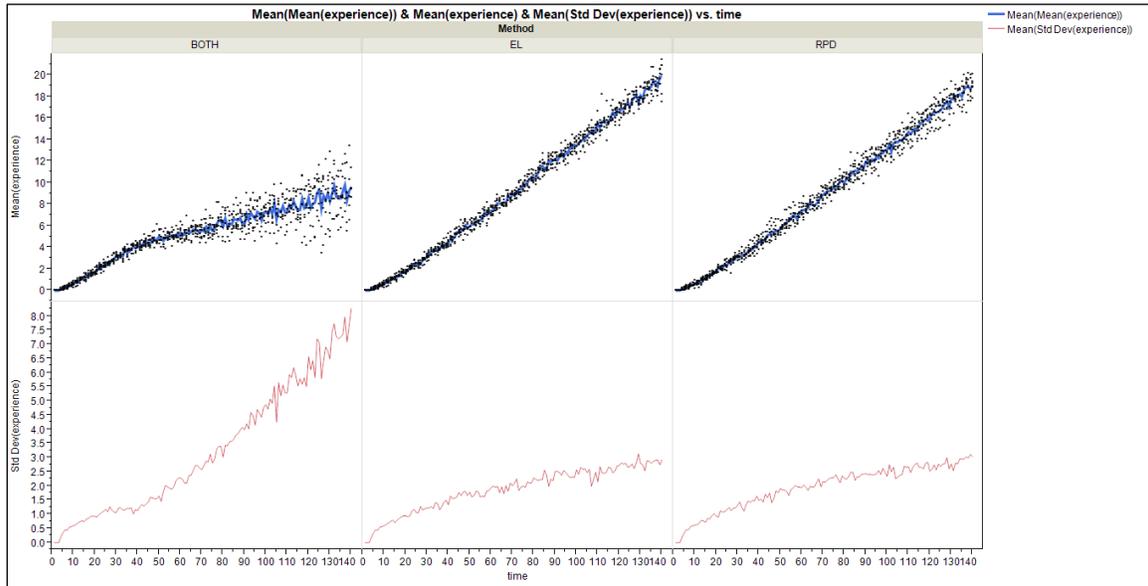


Figure 13. Entity Experience over Time.

3. Homophily and Communications

The degree of homophily was expected to differ between the I_P_U_S_Sp entities and the single Un_Pa_R_M_Ma entity. The earlier data indicating the slower response of the Un_Pa_R_M_Ma in terms of civil security issue stance pointed to the possibility that it was not receiving communications as readily due to its lower homophily link weigh with the other entities. The data shown in Figure 14 provides some evidence of this behavior, indicating that communications between I_P_U_S_Sp and Un_Pa_R_M_Ma averaged at an acceptance rate of 85.4%. In comparison, the communications between the I_P_U_S_Sp entities was accepted 86.1% of the time. More significantly, the volume of communications between I_P_U_S_Sp entites averaged 1.21 times a day, against 0.94 times a day for Un_Pa_R_M_Ma_1 to either of the other two entities. This indicated that the effect of homophily (determining the entities' desired to communicate with each other) was far more significant compared to trust (which determined acceptance of communications received). Comparison of the homophily link weights and trust levels between entities did not yield any other new findings.

		commDecision			
		RECEIVE	ACCEPT	RECEIVE	DONT_ACCEPT
sender	receiver	N	Row %	N	Row %
I_P_U_S_Sp_1	I_P_U_S_Sp_2	15261	86.55%	2371	13.45%
	Un_Pa_R_M_Ma_1	11323	85.57%	1909	14.43%
I_P_U_S_Sp_2	I_P_U_S_Sp_1	15261	85.67%	2553	14.33%
	Un_Pa_R_M_Ma_1	10899	85.26%	1885	14.74%
Un_Pa_R_M_Ma_1	I_P_U_S_Sp_1	10558	84.61%	1921	15.39%
	I_P_U_S_Sp_2	14608	85.98%	2382	14.02%

Figure 14. Communications Acceptance/Rejection Rates Between Entities in 3-Agent Scenario.

THIS PAGE INTENTIONALLY LEFT BLANK

IV. FURTHER TESTING AND EVALUATION

A. DESIGN PARAMETERS

The results and analysis of the initial set of design points suggested that the effects of decision method and trust were being overshadowed by other design factors in the model. The next phase of the testing and evaluation was thus developed to maximize the possible effects from these components of the cognitive architecture. In addition, factors that were found to be less significant or less relevant to test purposes were removed. The discount factor was fixed at 0.01, and only the casefiles based on 100 respondents were used.

The initial issue stance and OAB of entities was seen to have significant influence on the behavior and effect on scenario outcome. Several levels were tested, of which four were chosen for final set of design points. Most importantly, the periodic scripted action effect was removed and replaced with single action, as described in test scenario description in the next section. Table 6 shows the 24 design points that were used for the final run.

Design Point	Decision Method	Trust	Initial Stance	Design Point	Decision Method	Trust	Initial Stance		
951	RPD	ON	99% Adequate	963	RPD	ON	55% Adequate		
952		OFF		964		OFF			
953	EL	ON		965	EL	ON			
954		OFF		966		OFF			
955	BOTH	ON		967	BOTH	ON			
956		OFF		968		OFF			
957	RPD	ON		75% Adequate	969	RPD		ON	50% Adequate
958		OFF			970			OFF	
959	EL	ON	971		EL	ON			
960		OFF	972			OFF			
961	BOTH	ON	973		BOTH	ON			
962		OFF	974			OFF			

Table 6. Design Points for Final Run.

B. TEST SCENARIO

Six agents were utilized for the final round of testing. These comprised three I_P_U_S_Sp and three Un_Pa_R_M_Ma entities. The scenario was also expanded geographically – the two infrastructure nodes were placed at a distance of about 10 hex-grids apart, and the agents were distributed around them as shown in Figure 15. Each grid represents an area of approximately 1-mile radius.

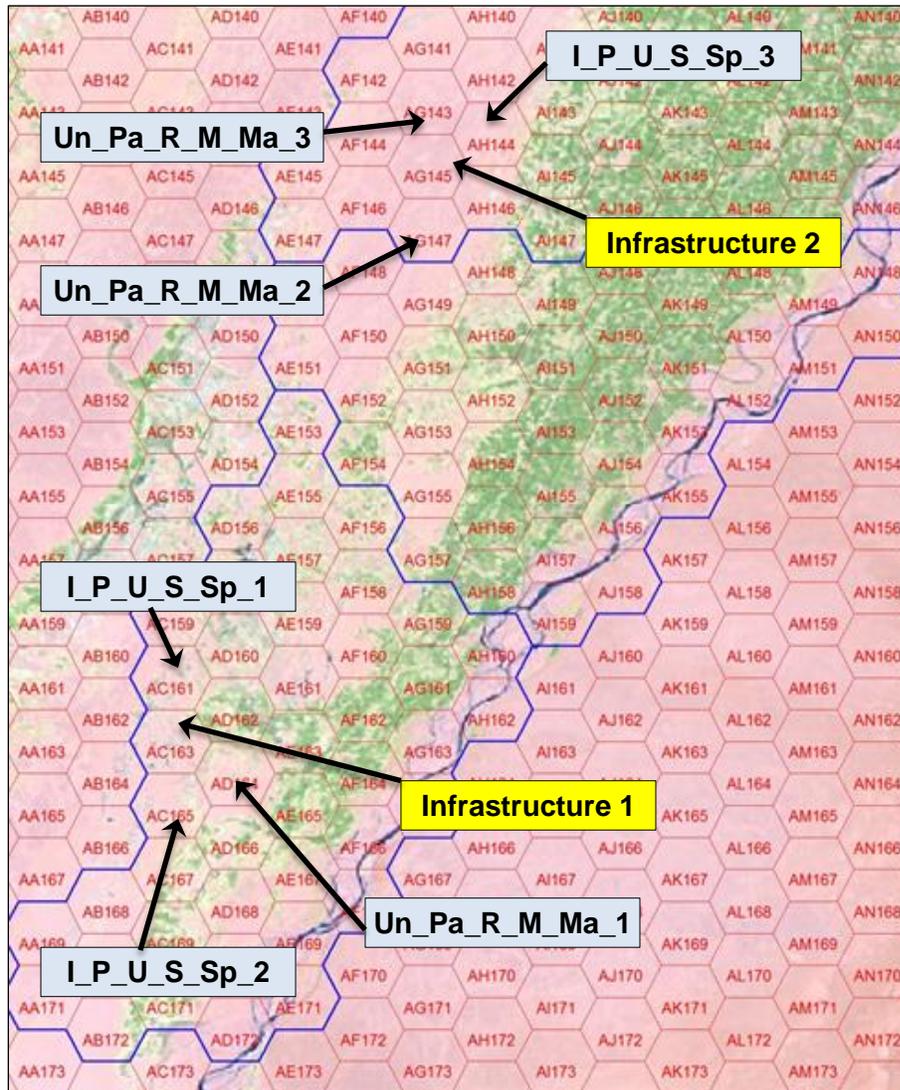


Figure 15. Map of Area of Operations (From Yamauchi, 2012).

With this set up, the effects of geographical location, communications between entities regarding infrastructure, and success rates of visiting the nodes will come into play. The effect of infrastructure visits was adjusted to have variable impact on entity stance—if an agent succeeds in restocking when he visits a node, there would be a 75% likelihood for a positive effect on stance, and a 25% otherwise. However, this is only one of the factors determining any overall change in stance, because the influence of other parameters also contributes to overall behavior choices and net change in issue stance.

The periodic scripted action used previously was replaced by a single action that occurred at a fixed time. The scenario was initialized with one of two infrastructure nodes *inoperable*, and the other at a *minimal* state (Table 7 provides the definition of infrastructure operation states). At day 90 of the scenario, the scripted action for CF to improve the inoperable infrastructure node takes place, restoring its state to *normal*. The operation state of the other node remains *minimal*. This setup causes entities to fail if they attempt to restock consumables from the first node prior to day 90, and to periodically fail when they attempt to restock from the second node throughout the scenario (essentially, only 1 of 7 attempts would succeed).

State	openTime	closeTime	numberServers	queueCapacity
Normal	360	0	1	10
Reduced	2	5	1	10
Minimal	1	6	1	10
Inoperable	-	-	-	-

Table 7. Definitions for Infrastructure Operation States.⁷

⁷ Several configurations for the initial state and state after scripted repair action were tested to develop this set of parameters and scenario settings, such as varying the queue capacity, transfer rates and resource capacity of the nodes. These settings mean that the node at *minimal* state will be available for 1 out of every 7 days. Entities attempting to restock on the days that it is closed will experience a failure in the action. Those visiting on the day it is open will most likely receive their requested resource, as the server and queue capacity is sufficient to provide for all entities in the scenario (unless balking or reneging occurs due to other entities being in the queue ahead of it). The *inoperable* state always fails to provide resource to the visiting entity.

Thus, the expected behavior is for entities to initially experience a decline in stance, due to the inability to receive the requested resource. Also, the choice of actions would favor Node 2 over Node 1. After the action of infrastructure improvement, Node 1 becomes more viable of the two, and agents who maintain exploratory behavior are expected to realize this, possibly communicate with other entities, and thereby cause action choices to shift in favor of Node 1. The effect on stance is expected to be favorable, since the entities would then experience a high success rate, and thus the overall scenario outcome should show an improvement of issue stance over time.

The scenario length for this set of tests was increased to 360 days, allowing for trends and outcomes to stabilize and possibly reach their steady state levels. Thirty replications were run for each design point, using the same seeds as before.

C. OUTPUTS

Additional dataloggers were used for this set of tests (Table 8), including new code that was added to the ongoing revisions of the CG Model. In particular, the BehaviorEffects-Datalogger was added to track all occurrences of entities visiting either infrastructure, and capture their success/failures as well as the resultant effect on their issue stance.

Parameter	Datalogger(s) Used	Description
Infrastructure Visits	BehaviorEffects-DataLogger	Record of infrastructure visits on both nodes, outcome (succeed / fail), and effect on civil security issue stance (increase / decrease / unaffected).
Other Parameters	Location-DataLogger State-DataLogger Behavior-DataLogger Action-DataLogger	Additional parameters were recorded for cross-referencing and checking purposes. These were the locations of entities (to check entity movement around the area), state of infrastructure nodes, behavior choices of entities and occurrence of scripted actions.

Table 8. Description of Additional Key Parameters Measured.

D. RESULTS

1. Civil Security Issue Stance

The effect of initial population stance on the scenario outcome is clearly visible in Figure 16. As expected, initial trend in civil security is negatively-sloped, given that the infrastructure in the scenario is unable to provide consumables for the entities most of the time. The introduction of the scripted event at Day 90 triggered the change in behavior, seen as either a reduction of the decline in issue stance, or a change in the direction of the trend.

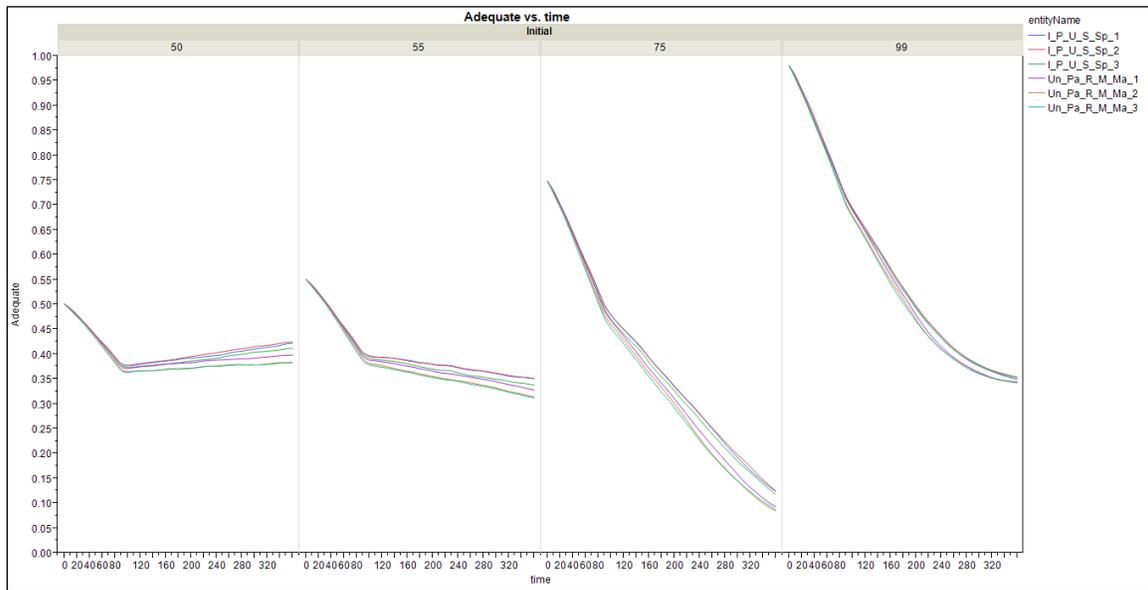


Figure 16. Civil Security Issue Stance for Different Initial Stance Levels.

In the CG Model, the initial issue stance determines the base effect from which the change caused by future actions are calculated. This implementation is responsible for the phenomena seen above, whereby the cases with a very high initial issue stance appears to be least affected by improvements brought about after the scripted action occurs. Further discussion of these effects is presented with the results of entity behavior and action selection in the next section.

Considering the case of 50% initial stance as an example (Figure 17), the decision method alone did not demonstrate significant effect on scenario initially. The trend of civil security issue stance over time for all entities followed a tightly bound range up till the point when the scripted action occurred. However, the effect of trust reduced the rate of change of entities' issue stances, resulting in a highly percentage of adequacy at the time the scripted action occurs. After day 90, the increase in choices available to the entities generated sufficient variation in the action-selection process to cause some degree of spread in the outcome at the end of the scenario as compared to the earlier simple scenarios. Figure 18 and Table 9 provide the breakdown of the civil security issue stance at the conclusion of the test scenario (day 360) for the 6 configurations of decision methods and trust. The results indicate that the overall scenario outcome is better (i.e., a higher percentage of the population feel that civil security is adequate) when the entities used both RPD and EL methods, compared to only one particular decision method.

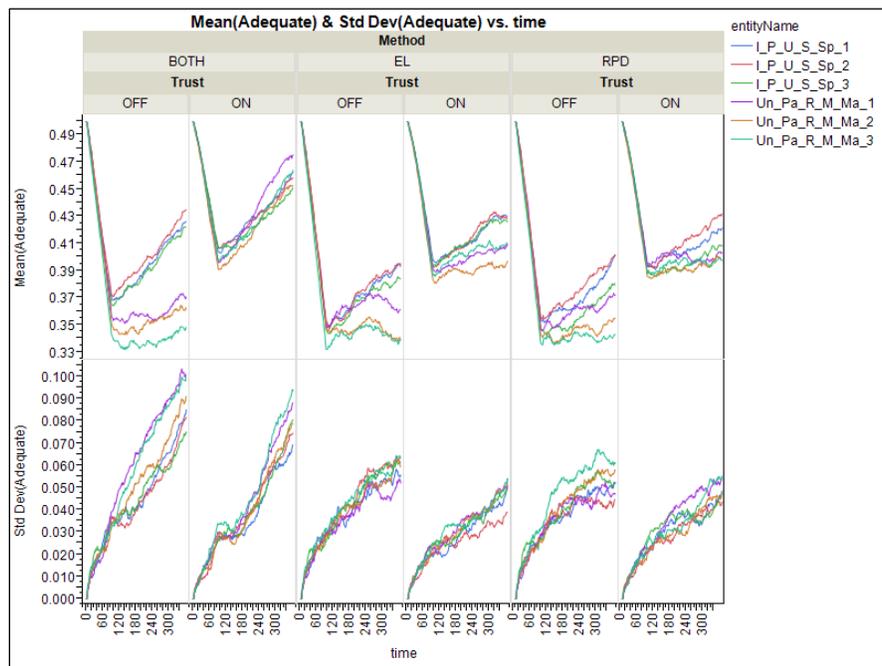


Figure 17. Civil Security Issue Stance for Initial 50% Adequate.

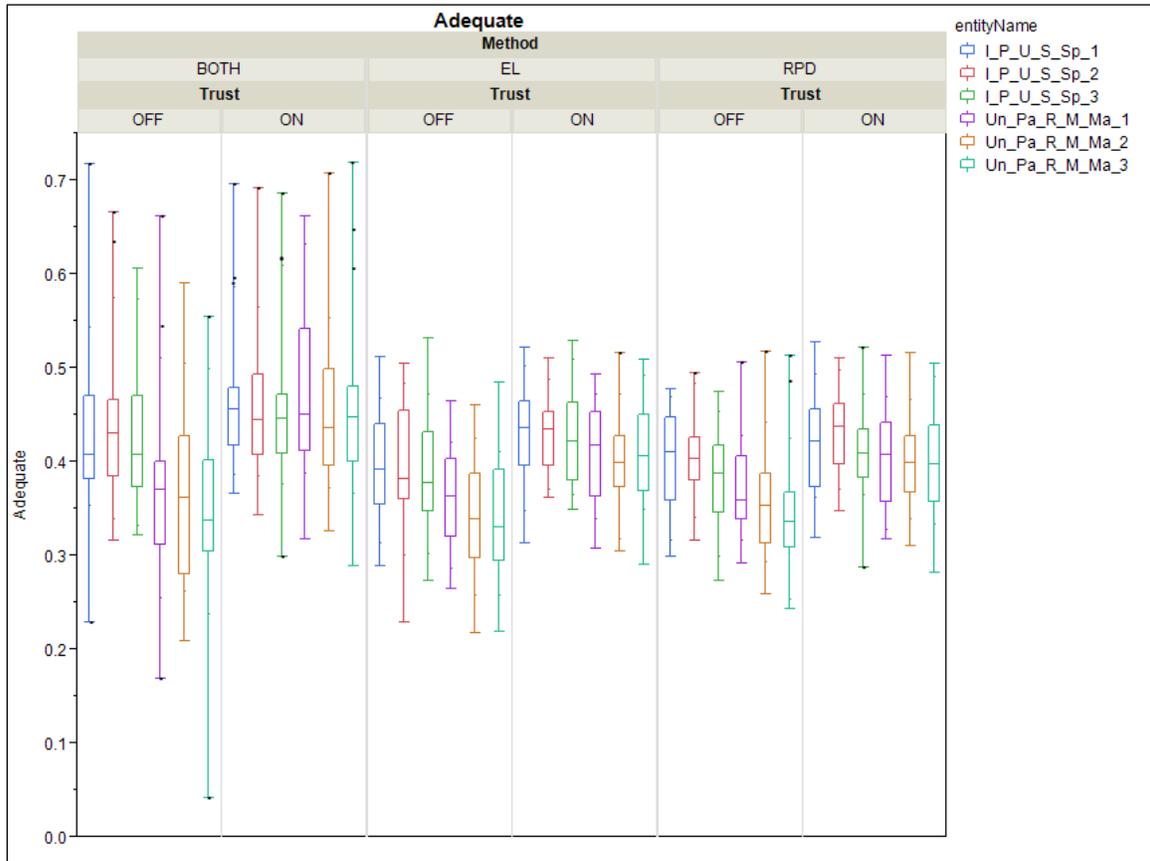


Figure 18. Distribution of Outcomes - Civil Security Stance at Day 360.

Configuration		Mean Stance (% Adequate)	Standard Deviation	95% Confidence Interval	
Method	Trust			Lower Bound	Upper Bound
BOTH	OFF	39.4%	9.5%	38.0%	40.8%
	ON	46.1%	8.1%	44.9%	47.3%
EL	OFF	36.9%	6.3%	36.0%	37.8%
	ON	41.7%	5.1%	41.0%	42.4%
RPD	OFF	37.6%	5.7%	36.8%	38.4%
	ON	41.0%	5.1%	40.3%	41.7%

Table 9. 95% Confidence Interval Levels of Civil Security Stance at Day 360 (Combined Mean across all Entities in Scenario).

2. Decision Method and Action Selection

The infrastructure-related choices made by entities in the final scenario provided further insight to their behavior and the effects of the decision methods. The actions selected and resultant effects are summarized in Figure 19, which includes the data from all 24 design points.

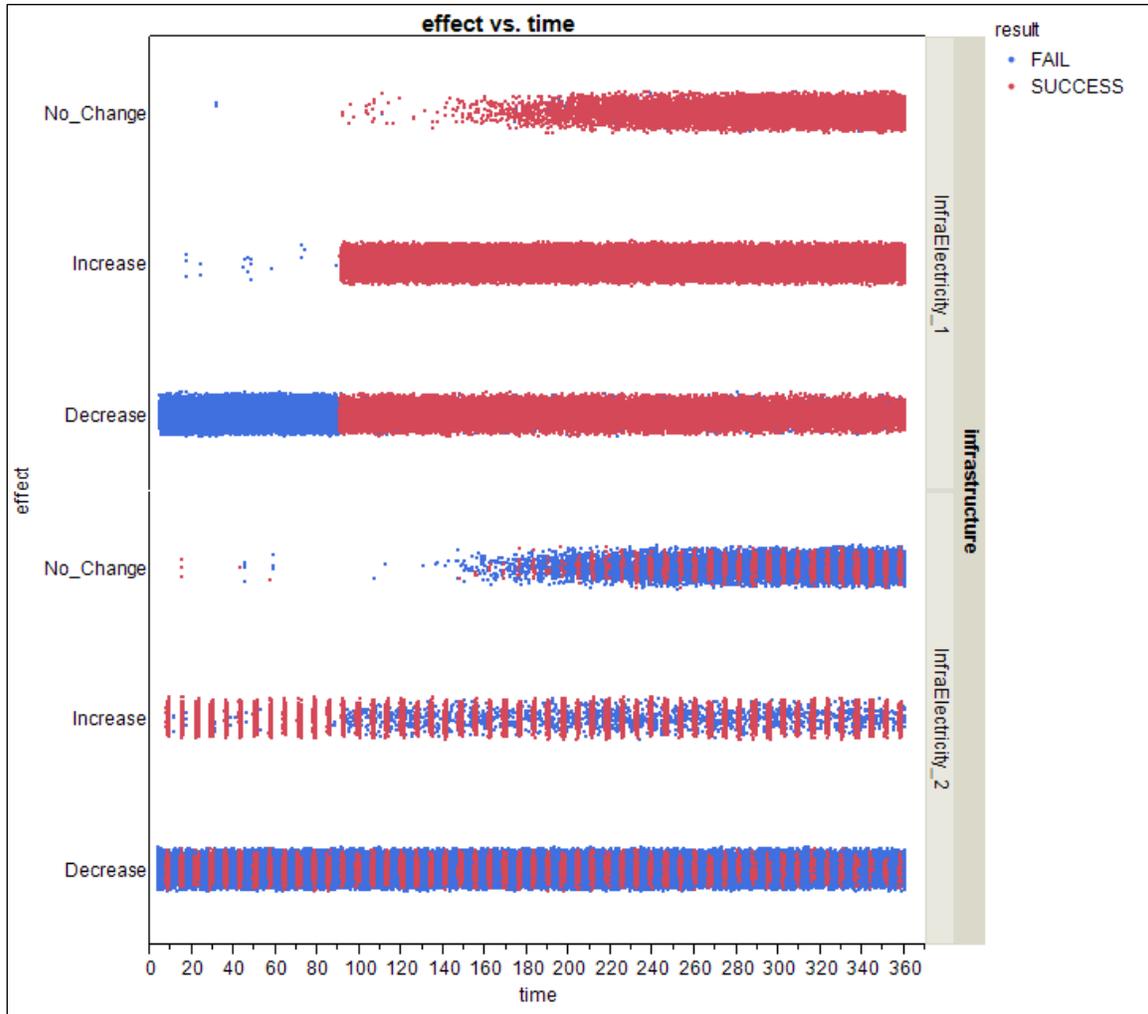


Figure 19. Infrastructure Node Visitation Outcomes and Effects.

The behavior of the entities provides a key insight that the outcome of an entity's visit to a node can generate both positive and negative effects on its issue stance, regardless success or failure to obtain the resource requested. In

particular, during the second half of scenario run time, there is a significant increase in instances of actions that do not cause any change to stance. The visitation rates of the two infrastructure nodes (Figure 20) provide a tell-tale sign that entity behavior is not ideal in the model / scenario—despite an total failure rate of 86.2% experienced with infrastructure node 2, entity behavior does not change to avoid it, as would be expected for a reinforced learner.

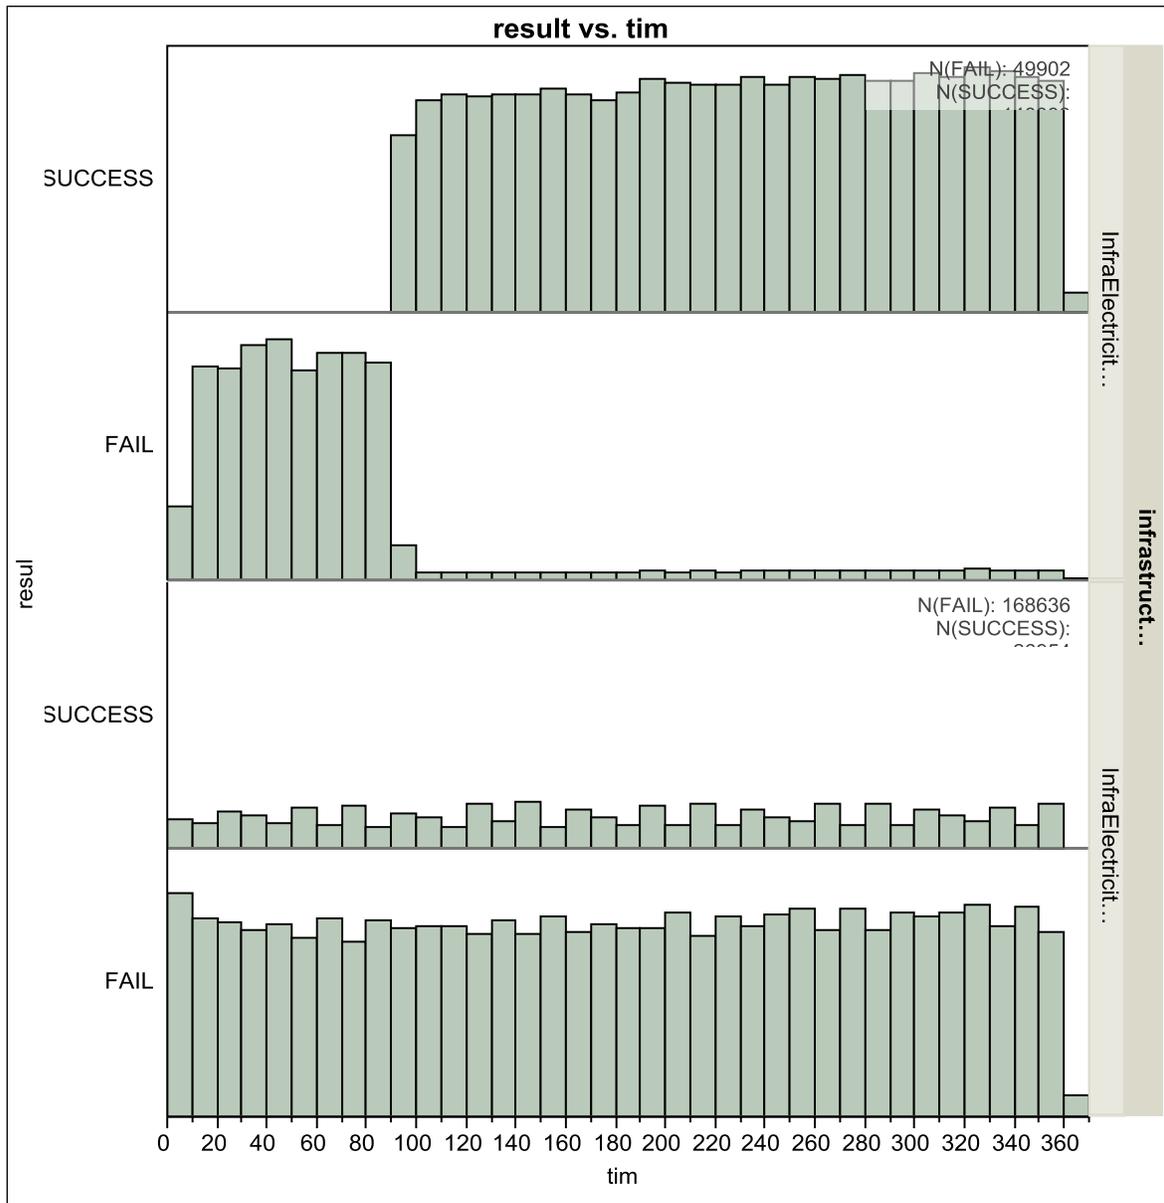


Figure 20. Infrastructure Node Visitation Rates and Outcomes.

Data from the action-selection process was used to investigate the cause of such agent behavior. Figure 21 plots the expected utilities of the three possible infrastructure-related actions on a logarithmic scale for all 24 design points in the scenario. The increase in expected utility of seeking a new provider corresponds to the occurrence of the scripted action at day 90; however, the action of remaining with an entity's existing provider also increases in value over time. This trend results in agent behavior that does not focus on either choice.

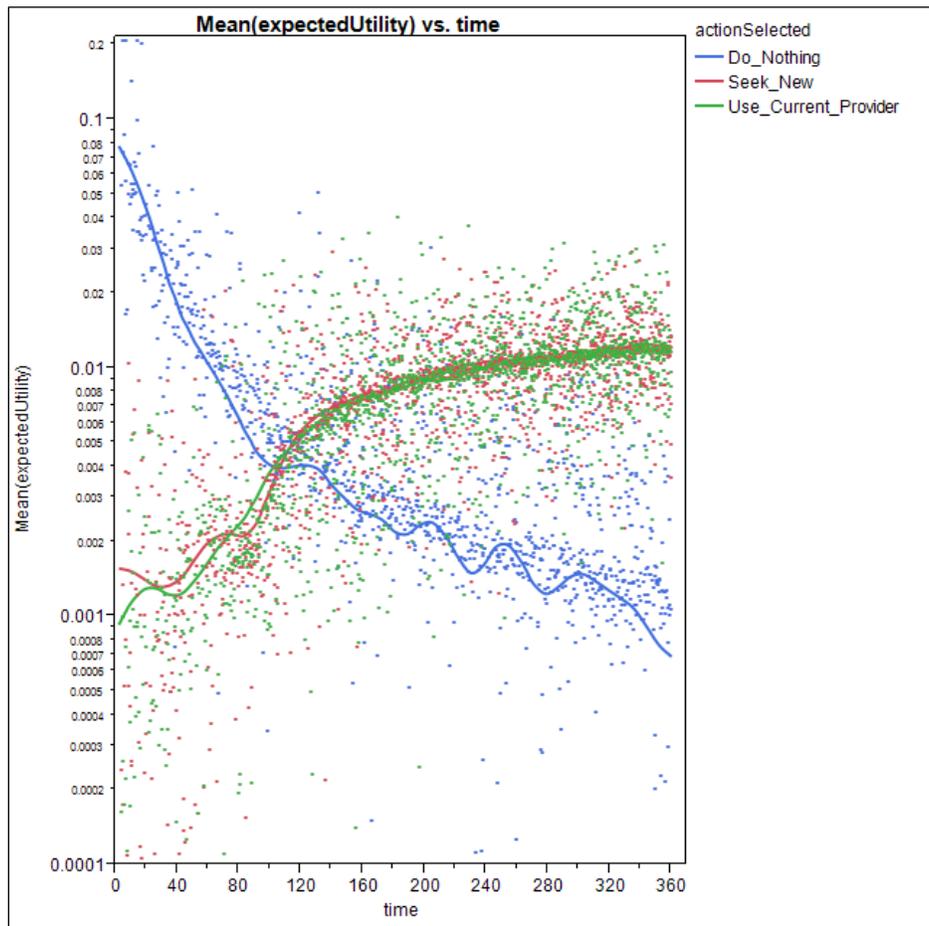


Figure 21. Expected Utility of Infrastructure-related Actions in 6-Agent Scenario.

Further analysis of the source code and consultation with the programmer (H. Yamauchi, personal communication, July 2012) revealed that the existing

algorithm for allocation of rewards to the actions does not account for the state of the entity, which explained the behavior observed in the infrastructure-related action selection process. Entities that visited a node and received an unfavorable outcome would have a higher probability of choosing to seek a new provider on their next action selection. However, upon switching to the better node, the expected utility for seeking a new node would be higher than the action of staying with that new provider. The resultant behavior would cause the agent to switch back and forth between nodes, seemingly with no regard to the outcomes from the infrastructure visits.

THIS PAGE INTENTIONALLY LEFT BLANK

V. CONCLUSION

The CG Model utilizes a highly complex cognitive architecture module in order to accurately and realistically depict the behavior of civilian populations in an IW environment. The critical process of entity decision making is based on well-accepted social science theories that provide a sound framework for the artificial intelligence of entities. The decision methods and trust module used in the CG Model were found to perform adequately, despite some deviations from expected behavior that were attributed to limitations in the implementation of these conceptual models.

A. EFFECTS OF DECISION METHOD

The process of decision method selection in the CG Model utilizes a reinforcement learning algorithm in two ways—as an exploratory approach, to allow entities to try out possible actions and build up their knowledge of expected utilities; and as a greedy approach, to simulate a RPD model of decision making. The test scenarios showed that the EL approach was adequate in generating agent behavior which performed as expected. The RPD approach generated similar scenario outcomes to the EL mode, in terms of overall trend and end state of civil security issue stance, behavior actions and interactions between entities. The combination of both methods, as implemented in the existing CG Model, generated scenario outcomes over a far larger range of possibilities, with close to twice as much variation as compared to either RPD or EL alone. However, the mean outcome was shown to be fairly similar across the design points tested. The effect of other parameters, in particular the initial stance of the entities, was far more significant in influencing the overall stance at the end of the scenario.

The significant increase in variance generated when both RPD and EL methods are used suggests that this implementation would be useful for the purpose of exploring potential outcomes for any given set of inputs, as it would cover a larger sample space.. However, continued development to independently

refine the RPD method would also be important to allow the model to better capture the effects of 'expert' entities (vis-à-vis a novice that would require several rounds of exploratory behavior to attain the same experience). Also, the existing cognitive architecture has limitations in associating utilities to state-action pairs instead of actions alone, which resulted in behavior that deviated from expectations, but still allowed entities to make choices and influence the outcome of the scenarios in a coherent manner.

B. EFFECTS OF TRUST

The inclusion of the trust module in the CG Model was shown to have a strong influence on the rate of change in issue stance of entities. This collaborates with the findings in Pollock's (2011) implementation; however, the outcomes of the test scenarios were shown to converge towards the same steady state regardless of the trust setting. The trust module thus serves as a buffer that delays the impact of actions in the area of operations, as its current form (as used in the test scenarios) only act to reject information. However, there is potential for it to influence scenario outcome, depending on the time frame allocated, and the frequency of actions occurring in the scenario.

C. OTHER FACTORS

The initial test scenarios demonstrated the strong impact that input parameters for a CG Model scenario can have. In line with the findings of earlier studies (Papadopoulos, 2010; Pollock, 2011), careful selection of these factors is crucial in order to build a realistic scenario that matches user requirements and expectations of agent behavior. The test cases showed, in particular, that the initial stance of the population was extremely significant.

D. TRACEABILITY OF ENTITY BEHAVIOR

The complexity of interactions in the CG Model makes tracing of entity behavior rather challenging. The process adopted in this study demonstrated the need to explore effects of different components of the CG at multiple levels,

ranging from the isolation of single factors to larger scenarios with multiple parameters being evaluated. The dataloggers built into the existing CG Model served as valuable tool for recording the immense amount of data generated in each replication and design point.

The experimentation done in this thesis has assisted the ongoing development of the CG Model. Several revisions of the code were made to adjust settings and rectify minor anomalies in the entity behaviors. The creation of new dataloggers by TRAC-MTRY programmers would also provide for future testing and evaluation efforts, and improve the traceability of entity behavior.

E. FUTURE WORK AND RECOMMENDATIONS

The analysis of the effects of decision methods in the CG Model revealed a few aspects of the cognitive architecture that could be improved. The greedy reinforcement learning approach used for the RPD method and the limitation on state-action pair association in the EL method are two key areas that could be investigated for future developments.

In terms of analysis and testing of the cognitive architecture, several areas have been identified that could benefit from further study:

1. The test scenarios used in this study utilized only two entity prototypes, which posed a constraint on the extent of differences in homophily and possible interactions between them. Expansion of the scenario to include more agent types would serve to test the effect of homophily and communications to a greater extent.

2. The EL method is applicable to a wide range of actions that entities could undertake in the CG Model. The testing of infrastructure-related actions in this study was limited by the lack of accounting for entities' existing states (current resource provider). Testing of the EL method in other contexts, in

particular for scenarios or actions that are less/not dependent on state would serve to build up further understanding of the action selection process in the CG Model.

3. The current implementation of trust in the CG Model acts to restrict information flow to an entity. An opposite effect could be modeled such that an entity receiving percepts from a highly trusted counterpart would be influenced to a greater extent than normal. This would allow shifts in scenario outcomes in either direction as a result of trust, instead of the single-direction “buffering” effect that was observed in this study. However, such an implementation would increase the complexity of the CG Model even further.

This study has shown that the decision methods and trust module in the cognitive architecture are significant components in the CG Model. However, their effects are not always visible in terms of measurable outcomes such as issue stance of entities and overall trends in agent behavior. The test scenarios involved simplistic settings and did not exhibit any degradation of performance (e.g., computation / simulation time). However, with full-scale wargaming scenarios, the removal or deactivation of some components may become an acceptable tradeoff.

LIST OF REFERENCES

- Ajzen, I. (1985). From intentions to actions: A theory of planned behavior. In J. Kuhl & J. Beckmann (Eds.), *Action-control: From cognition to behavior* (pp. 11–39). Heidelberg, Germany: Springer.
- Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50, 179–211.
- Alt, J. K. (2010). Agent frameworks for discrete event social simulation. *Proceedings of the 19th Conference on Behavior Representation in Modeling and Simulation*. Charleston, SC.
- Alt, J. K., Jackson, L. A., Hudak, D., & Lieberman, S. (2009). The cultural geography model: Evaluating the impact of tactical operational outcomes on a civilian population in an irregular warfare environment. *The Journal of Defense Modeling and Simulation: Applications, Methodology, Technology*, 6(185), 185–199.
- Baez, F., Alt, J., Brown, R., Lieberman, S., Ehlschlaeger, C., Zwicky, W. R., ... Yamauchi, H. (2010). *Cultural Geography (CG) Model Enhancements for Irregular Warfare (IW) Analysis*. Monterey, CA: TRAC-MTRY.
- Buss, A. (2011). *Discrete Event Simulation Modeling* (Technical report). MOVES Institute, Naval Postgraduate School.
- Department of Defense. (2008). *Irregular warfare* (DoD Directive 3000.07). Retrieved from <http://www.dtic.mil/whs/directives/corres/pdf/300007p.pdf>
- Ferris, T. P. (2008). *Modeling methodologies for representing urban cultural geography in stability operations* (Master's thesis). Naval Postgraduate School.
- Fisher, W. R. (1984). Narration as a human communication paradigm: The case of public moral argument. *Communication Monographs*, 51(1), 1–22.
- Hudak, D. & Baez, F. (n.d.). *Cultural geography modeling and analysis in helmand province*. Technical Report. Monterey: TRAC-MTRY.
- Joint Chiefs of Staff. (1995). *Joint doctrine for military operations other than war* (Joint Publication 3–07). Washington, D.C.: Joint Staff Pentagon
- Klein, G. (1999). *Sources of power: How people make decisions*. Cambridge, MA: MIT Press.

- Klein, G. A., Calderwood, R. & Clinton-Cirocco, A. (1986). Rapid decision making on the fire ground. *Proceedings of the Human Factors Society 30th Annual Meeting*.
- Klein, G. A. (1998). A recognition-primed decision model of rapid decision making. In Klein, G.A., *Recognition-primed decisions* (pp. 138–147). Retrieved from Defense Technical Information Center OAI-PMH Repository.
- McPherson, M., Smith-Lovin, L. & Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27, 415–444.
- Modeling & Simulation Coordination Office [MSCO]. (n.d.). Home Page of the Department of Defense Modeling & Simulation Coordination Office. Retrieved from <http://www.msco.mil/descMSCO.html>
- Ng, E. H. (2012). Keynote address by Minister for Defence Dr Ng Eng Hen at the *2012 Singapore Armed Forces Learning Symposium*. Singapore.
- Ozcan, O. (2011). *Balancing exploration and exploitation in agent learning*. (Master's thesis). Monterey, California: MOVES Institute, Naval Postgraduate School.
- Ozcan, O., Alt, J., & Darken, C. J. (2011). Balancing exploration and exploitation in reinforcement learning. *Proceedings of the Twenty-Fourth International Florida Artificial Intelligence Research Society Conference*. FL.
- Papadopoulos, S. (2010). *Reinforcement learning: A new approach for the cultural geography model* (Master's thesis). Monterey, California: MOVES Institute, Naval Postgraduate School.
- Perkins, T. K., Pearman, G. M., Baez, F. (n.d.). *Representing civilian population interaction with infrastructure and essential services in stability operations*. Monterey, CA: TRAC-MTRY.
- Pew, R. W., & Mavor, A. S. (Eds.). (1998). *Modeling human and organizational behavior: Application to military simulations*. Washington, D.C.: National Academy Press.
- Pollock, S. (2011). *Representing trust in cognitive social simulations* (Master's thesis). Monterey, California: MOVES Institute, Naval Postgraduate School.
- Proctor, R. W. & Zandt, T. V. (2008). *Human factors in simple and complex systems* (2nd ed.). Boca Raton, FL: CRC Press.

- Russell, S. & Norvig, P. (2010). *Artificial intelligence: A modern approach* (3rd ed.). Upper Saddle River, NJ: Pearson Education, Inc.
- Seitz, T. (2008). *Representing urban cultural geography in stabilization operations: Analysis of a social network representation in pythagoras* (Master's thesis). Monterey, California: Naval Postgraduate School.
- TRADOC Analysis Center – Monterey [TRAC-MTRY]. (2009). Irregular warfare research to represent civilian populations in stability operations [PowerPoint Slides]. Retrieved from http://www.mors.org/UserFiles/file/meetings/09iw/pres/WG2_Jackson_J.pdf
- U.S. Army PEO STRI. (2012). *One Semi-Automated Forces Project Office Description*. Retrieved from <http://www.peostri.army.mil/PRODUCTS/ONESAF/>
- Watkins, C. J. C. H. & Dayan, P. (1992). Technical note: Q-Learning. *Machine Learning*, 8, 279–292.
- Watkins, C. J. C. H. (1989). *Learning from delayed rewards* (Doctoral dissertation). Retrieved from <http://www.cs.rhul.ac.uk/~chrisw/>
- Wiedemann, M. (2010). *Robust parameter design for agent-based simulation models with application in a cultural geography model* (Master's Thesis). Monterey, California: Naval Postgraduate School.
- Yamauchi, H. (2012). *Cultural geography model (version 0.9x)*. Draft Technical Report. Monterey: TRAC-MTRY.

THIS PAGE INTENTIONALLY LEFT BLANK

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California
3. Director, TDSI
Temasek Defence Systems Institute
Singapore, Republic of Singapore
4. Senior Manager, TDSI
Temasek Defence Systems Institute
Singapore, Republic of Singapore
5. COL Lim Soon Chia
Deputy Chief Research & Technology Officer
Defence Research & Technology Office, Ministry of Defence
Singapore, Republic of Singapore