A Demonstration of the PMF-Extraction Approach: Modeling The Effects of Sound on Crowd Behavior

Jason B. Cornwell[1], *Barry G. Silverman*[1,2], *Kevin O'Brien*[1], *Mike Johns*[3]

1 - Ackoff Center for the Advancement of the Systems Approach (ACASA),
2 - Systems Engineering Department and Institute for Research in Cognitive Science
3 - Center for Human Modeling and Simulation, CIS Department

University of Pennsylvania,
Philadelphia, PA 19104-6315
jasonc@seas.upenn.edu
barryg@seas.upenn.edu
eon@seas.upenn.edu
mjohns@seas.upenn.edu

Key Words:
Human Behavior Model, Performance Moderator Function, Sound, Noise, Music

ABSTRACT: *The vast majority of psychology, sociology, and other social-science literature describing human behavior and performance does not reach the eyes of those of us working in the modeling and simulation community. Our recent work has been concerned with the extraction and implementation of Human Behavior Models(HBMs)/ Performance Moderator Functions(PMFs) from this literature. This paper demonstrates how our methodology was applied to extract models of the effects of music and sound on both individuals and groups and to implement them in a simulated environment. PMFs describing how several classes of sound affect decision-making and performance were constructed based on well-established psychological models. These PMFs were implemented in a simulation of protesters and security guards outside a prison that demonstrates how the presence of chanting and music changes the response of protesters to police aggression. The extraction of PMFs from the literature, the synthesis of a coherent, cohesive model, and the implementation and results of the simulation are discussed.*

1. Introduction

The developers of models and simulations involving human behavior and performance are not typically psychologists or sociologists. However, in order to craft realistic and plausible simulations for training and analysis developers should draw from the massive body of human performance models and data available in the social science literature. The benefits of using pre-existing models from psychology and sociology – as opposed to creating ones own – are both numerous and obvious. Theoretical models can be selected that are robust, well respected, and have been empirically tested. Moreover, these models can be identified and implemented in existing systems relatively quickly. Our present line of research aims to demonstrate the efficacy of this approach.

There are a staggering number of human behavior models (HBMs) and performance moderator functions (PMFs) detailed in social science journals. An HBM/PMF captures a dose-response type of relationship between a performance moderator and the level of performance. These moderators reflect significant dimensions of individual and group differences (e.g., intelligence, skill, judgment, leadership, emotion, organizational culture, motivation, dedication, slips/lapses/biases) as well as external stressors on individuals and/or groups (e.g., task time, noise, fatigue, stress, opponent actions, etc.). HBM/PMFs are of variable validity and relevance, so a good deal of effort is required to sort though the literature and catalog useful, valid, and relevant models. Much of our recent work has involved the collection and categorization of anthologies of HBM/PMFs [1]. To date we have collected several hundred, all of which have been condensed into structured abstracts and rated based on their validity to facilitate rapid implementation [2]. This collection amounts to a tiny fraction of the potentially useful HBM/PMFs available. The collection and categorization of all such HBM/PMFs is not a reasonably achievable goal, so we have focused our attention on the compilation of limited anthologies and the development of our extraction and implementation methodology.

Another facet of our research has been the development of a general cognitive architecture in which to deploy HBM/PMFs. Our architecture allows for a wide and flexible set of behaviors and representations and, although we have built a limited initial simulation test bed for it, it is designed to be portable to other simulation environments. Our goal is to give simulation developers a tool that lets them quickly and easily either select from a wide range of pre-catalogued HBM/PMFs or cull their own from the literature, drop them into a general cognitive architecture, and run this architecture within their existing simulations. The basic architecture is described briefly below and is covered more fully elsewhere in these proceedings [3].

To demonstrate our approach of HBM/PMF extraction and implementation within a general architecture, we chose to model the effects of sound on the behavior of both individuals and crowds. This choice was practically motivated. We had previously designed a series of scenarios within our simulation test-bed and cognitive architecture that explored "crowd equilibrium tipping" events and the conditions under which rioting can occur [4]. Sound had not been modeled in these scenarios and we expected that its inclusion would be a marked improvement in the validity of the simulation.

2. Sound Literature and PMFs

The effect of sound on behavior is too large and complicated an issue to be tackled in its entirety, so we broke it up into a series of smaller components that represent the specific aspects of sound that we were interested in considering for inclusion in the simulation: noise, music, and event-specific sound. Event-specific sound includes those sounds that are causally inseparable from the event that created the sound such that the behavioral response to the sound itself is subsumed by the response to the event. For example, it makes little sense to consider the effect of the sound of an explosion independently from the effect of the explosion itself. For our purposes music includes any rhythmic individual or group expression (drumming, chanting, etc). Noise encompasses all those sounds that are neither explicitly musical in nature nor overshadowed by the event that produced them.

These divisions are reflected in the separate bodies of literature that deal with each sub-topic. Noise has been exhaustively studied by psychologists, engineers, and urban planners, while the majority of the studies concerning the effects of music on behavior come from the music therapy community. Research on event-specific sound is, unsurprisingly, distributed across a variety of domains. Because event-specific sound is not a research area in and of itself, and because we can implement desired event-specific sound effects within our existing simulation without additional modifications (see below), we chose not to extract event-specific sound PMFs and instead focused our energies on the noise and music PMFs.

We extracted our general noise PMF from Broadbent's [5] excellent review of the effects of noise on human performance. Broadbent's principal conclusion, and the basis for our PMF, is the idea that as noise increases arousal increases. Broadbent uses this hypothesis to explain effects demonstrated in a wide variety of experiments running the gamut from measures of general cognitive ability under differing amounts of noise to signal detection to the performance of factory workers exposed to varying amounts of noise. Our own survey of the literature supports Broadbent's conclusion. A wide variety of phenomena can be explained in terms of general arousal due to noise: aggressive tendencies under noisy conditions [6], anxiety in noisy social situations [7], and even differential evaluations of group dissenters and conformists under noisy conditions [8] easily fall within Broadbent's framework. Given the simplicity of Broadbent's explanation and the consistency with which it explains a wide variety of reported phenomena, developing a PMF for the effects of general noise was not dificult. Our final PMF, which we consider to be quite valid, states that activation, arousal, and/or stress (depending on the simulation) have a positive linear correlation with general noise level.

$$A_t = f_d(A_{t-1}) + k \cdot n_t$$

A_t: total activation (and/or arousal, stress) at time t
f_d: decay function
k: constant representing susceptibility to noise
n_t: noise at time t

Developing a PMF for the effects of music was less straightforward. Although much research has been conducted in this area, there are no literature reviews that compare to Broadbent's noise survey in both quality and relevance. However, in our investigation of the music therapy literature, we found clear and consistent themes around which we constructed our PMF. The following papers provide especially compelling examples: Cassity [9] demonstrated that psychiatric patients who participated in group musical activity yielded significantly improved peer acceptance and group cohesiveness ratings as compared to patients who participated in non-musical activity. Anshel and Kipper [10] reported that participants in group-singing exercises exhibited a marked increase in intra-group trust and cooperation. Galizio and Hendrick [11] showed that political messages delivered musically have a greater propensity to change opinions than do political messages delivered via speech alone.

These studies tell a consistent story. When producing music, groups of people tend to be more single minded in the pursuit of their goals and individuals feel a stronger bond with the rest of the group than under normal conditions. At the level of the individual this may be interpreted as a reduction in physiological and safety concerns and an increase in the emphasis placed upon belonging and furthering the objectives of the group at the expense of other personal interests.

We did not find sufficient quantitative predictions in the literature to construct a moderator function that describes this finding analytically. Even if we had, any PMF that we might have pieced together that described the phenomenon in question – that musical performance yields a general shift in values away from the personal and towards the social – could not have been both meaningful and simulation-independent. We therefore left the PMF in written form and constructed a mathematical model when it came time to implement the PMFs within our architecture. We believe that the PMF we extracted from the music therapy literature is valid, but it is extremely general and not grounded in quantitative predictions and so is considerably less robust than the noise PMF.

As stated before, we chose not to focus on the implementation of event-specific sound effects in our simulation, although they do exist and are implicit in certain situations, as will be discussed in section 4 below. Our focus is on the development of general simulation architectures and frameworks. The implementation of individual event-specific sound effects would not help us towards that goal.

3. Simulation Architecture

Before we discuss the implementation of our sound PMFs, a brief overview of our simulation design and agent mind architecture is needed. Silverman [1] described the general topology of our Java-based agent but it has gone through several subsequent revisions and is significantly more developed. (A much more detailed explanation of our architecture is presented elsewhere in these proceedings [3].)

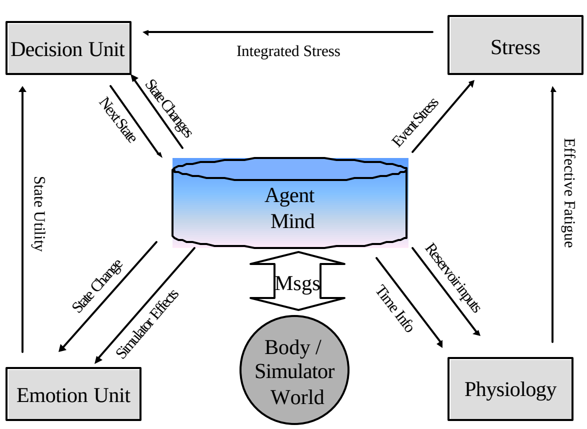

Figure 3.1: Agent Architecture

Our agent mind is built around four interconnected, interchangeable components: a physiology unit, stress unit, emotion unit, and decision unit. These components communicate with each other via messages that are compatible with Agent Communication Language so that they can be readily swapped with replacement components or left out of the simulation altogether. Our agent is, therefore, a multi-agent system with individual sub-units that can be thought of as agents in their own right. Figure 3.1 depicts the connections between the different units.

The decision unit processes internal markov chains representing all of the possible states of the agent. A simple chain taken from a civilian agent in our sample scenario is sown in Figure 3.2. The shaded circles are reactive states. Our agent can be automatically bumped into one of these states as a result of events in the simulation. For example, if a security agent attempts to arrest a protesting agent, that protesting agent will be bumped automatically into the "Deal with Arrest" state and must choose how to proceed from there. When a decision is called for, the decision unit sends queries to the emotion unit to request the expected utility of each sequence of steps available to the agent and calculates a plan that maximizes that utility. This decision process is further constrained by the stress unit, as described below.

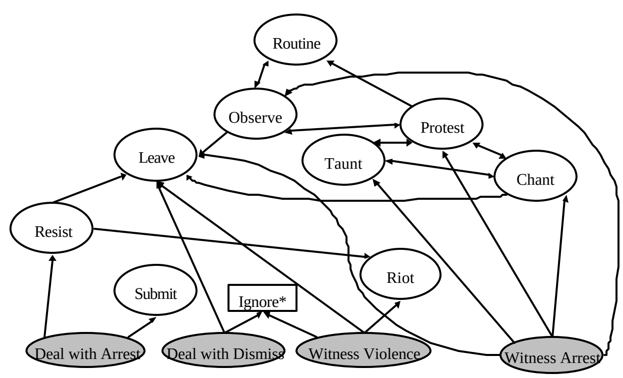

Figure 3.2: Simple Markov Chain

The emotion unit contains hierarchical, dynamic representations of the agent's goals, preferences, and cultural biases. Each of these factors is taken into account in the utility calculations as described by Ortony, Clore, and Collins [12]. The possible states of the world are attached as leaves to the value hierarchies, which are tree data structures representing the agent's goals, standards, and preferences. The skeleton of one such goal structure is presented in Figure 3.3. The structures are designed such that multiplying up the hierarchies from a leaf node – or possible world state – yields utility values for the agent for that state. Figure 3.3 leaves out both the leaf nodes and the values for all but the top links in order to be more readable. The values associated with links extending down from any given node should sum to 1, and indicate the relative importance the agent currently places upon each sub-tree. A new feature of our architecture allows the simulation to send messages to the emotion unit to alter the contents of the hierarchies. For example, an event in which another agent is killed might trigger a message that adjusts the agent's top-level value hierarchy nodes such that the value of the links to Physiology and Safety increase while the links to Esteem, Belonging, and Actualization decrease. See Johns et. al. [13] for a complete discussion of the emotion unit.

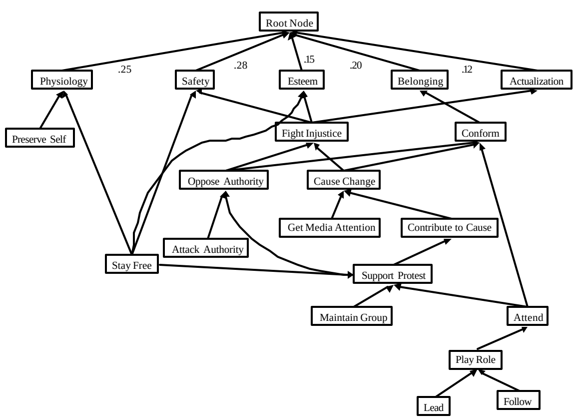

Figure 3.3: Sample Value Hierarchy

The stress unit further constrains the options available to the decision unit. Based on the Janis-Mann integrated stress model [14], the stress unit tracks the agent's overall arousal, or integrated stress. At very low stress levels, the decision unit is turned off and the agent absentmindedly continues to execute its current plan without evaluating new options. At very high stress levels, the decision unit is forced to evaluate a restricted set of options and looks only one step ahead, leading to panic and hasty decisions in novices and recognition primed decision-making in experts. Only when the agent's level of activation is within an intermediate range is the decision unit allowed to choose optimally. The stress unit maintains reservoirs that monitor event stress, fatigue, and time pressure, all of which are affected by events in the simulation. The integrated stress value is calculated based on the status of these reservoirs.

Lastly, the physiology unit maintains a set of physiology reservoirs representing fatigue, hunger, sickness, etc. These reservoirs are used to fill the fatigue reservoir in the stress unit.

Several other mechanisms complete the agent. Message handlers can be customized to modify any aspect of the agent. They can, for example, bump an agent into a particular state in the decision unit or change value hierarchies in the emotion unit. We also introduced "seed states" into the agents' markov chains. These states are disconnected from the rest of the chain until some event in the world introduces a link between the seed state and another state in the chain. This allows, for example, one agent to observe and learn a behavior from another agent or, more accurately, to realize that a particular option is available because the agent saw another agent performing that action.

4. PMF Implementation

Our simulation architecture is quite general, but it should be flexible and extensible enough to allow for a wide range of specific behaviors. One motivation for implementing sound PMFs into our architecture has been to test this extensibility. Continually adding specific PMFs shows us those aspects of our agents that are sufficiently flexible and exposes those that are insufficiently developed.

The noise PMF we adapted from Broadbent [5] fit quite readily within the existing architecture. We created a noise reservoir within the physiology unit that implemented the PMF as follows: Noisy events in the environment broadcast messages that describe their volume. Agents within range that receive the message broadcasts send a message to their physiology unit instructing it to update their noise reservoir. The noise reservoir, which decays over time without input, feeds into the integrated stress value that, in turn, constrains the decision unit. Our simulated noise therefore has the effect of increasing anxiety and limiting attentional capacity – exactly the pattern Broadbent [5] presented.

This noise implementation provides the capacity for ready implementations of particular event-specific noise effects as well, and also underlies the music PMF. Any event that produces sound will add to the agents' integrated stress values. Other, more specific effects can be implemented on top of this basic mechanism.

Our music PMF did not fit into the architecture quite as seamlessly, as we were unable to generate a quantitative PMF. The behavioral description calls for a shift in values from individual goals towards group goals as well as an increased feeling of connectedness and belonging to the group. Implementation of these temporary shifts in values can only take place within the emotion unit's value hierarchy. In previous versions of our architecture these values had been static, but it was immediately obvious that a large set of behaviors would only be possible given value hierarchies that could fluctuate over time. To implement the music PMF, we modified the value hierarchies so that the weight of each link between nodes could be changed via a message, but would then decay back to its original value over time. Refer back to Figure 3.3 for a sample value hierarchy of an agent used in our simulation demo. When an agent is involved in musical activity, it sends a message to its emotion unit that modifies the weights of the top-level value nodes. Esteem, Actualization, and Belonging become more important to the agent, while physical comfort and safety diminish in importance. When the music stops, these values decay back to their default settings.

5. Sample Scenario Design and Results

We designed a sample scenario in order to deploy and test our sound PMFs. This scenario was built upon a central question – how does chanting affect the interactions between protesters and security forces at a protest? To address this question, we constructed two simulations that differed in one respect only: In the first simulation, protesters did not have access to chanting as a possible state on their markov chain. In the second simulation, one protester had the chanting state attached to its markov chain and could "teach" other agents how to chant via seed states.

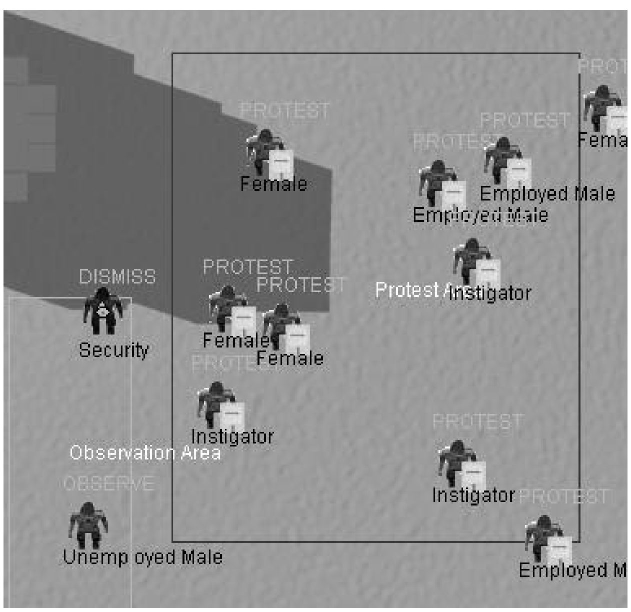

Figure 5.1: S1-Protesting despite dismissal

The simulations take place outside of a prison where a group of protesters are facing a lone security guard. This guard's value hierarchy was designed to make him quite aggressive and gave him a tendency to violently arrest protesters that do not disperse when asked. The protesters themselves were drawn from several different agent types. Unemployed male, employed male, female, and provocateur agents were all represented in the group. We set the agents' initial states and positions in the simulation such that the security guard would try to disperse the protesters at the start of the scenario and that the protesters would be disinclined to leave the scene as ordered, as their value hierarchies led them all to favor staying and protesting, as is shown in figure 5.1.

The first simulation was relatively short. The protesters ignored the guard's dismissal, which drove the guard to try to arrest one of the protesters (an instigator) who was taunting him. The protester resisted arrest, so the guard chose to attack him in response. The other protesters, whose value hierarchies had been shifted towards self-preservation by the guard's violent attack, quickly dispersed and scattered all over the map as is shown in Figure 5.2.

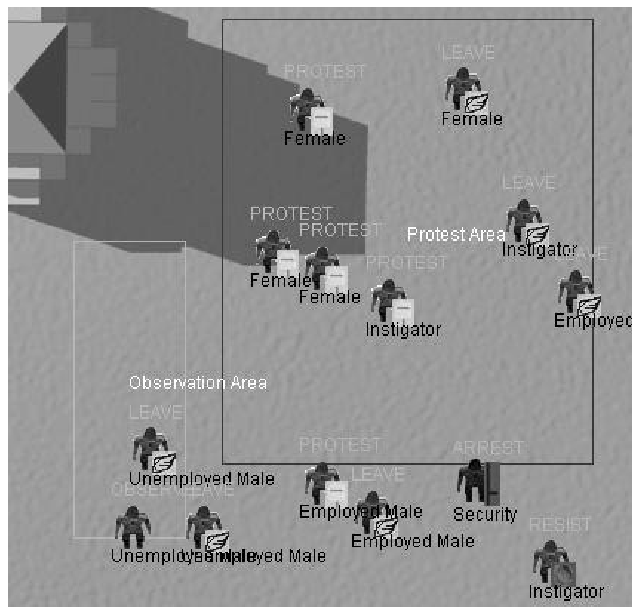

Figure 5.2: S1-Protesters disperse

The second scenario turned out quite differently, however. When the guard attempted to disperse the crowd, the protester who knew how to chant did so, and soon all of the protesters decided to chant. Rather than submit to the guard and leave the scene out of self-interest, the chanting shifted the protesters' value hierarchies away from self-preservation and emphasized esteem, belonging and actualization. As a result, the protesters found greater utility in standing their ground to support the common cause. The noise of the chanting combined with its failure to achieve its own goals raised the integrated stress level of the guard, who panicked and attacked a protester. Witnessing the attack, the provocateur agents seized the opportunity and began a riot. Exhausted, the guard retreated. Figure 5.3 depicts this senario.

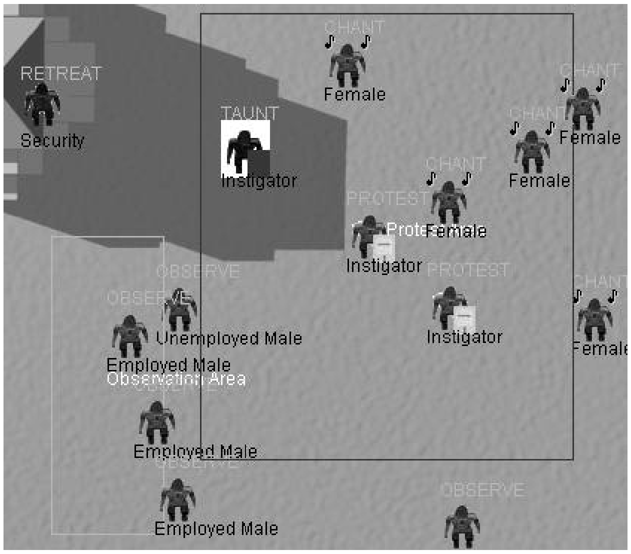

Figure 5.3: S2-Chanters Stand Ground

6. Discussion

This paper has outlined a general methodology for the extraction of PMFs from social science literature, the implementation of those PMFs within a general cognitive architecture, and the development of scenarios that both test the capabilities of the PMFs and expand the capabilities of the cognitive architecture in the process. We have demonstrated that our architecture is both robust and extensible enough to support the relatively abstract social concepts manipulated in the music PMF with only a few modifications. The noise PMF required even less effort to integrate into the architecture.

We have also shown that incrementally adding individual PMFs exposes the limitations of our architecture and allows us redesign pieces of it without reengineering the simulation in its entirety. Our initial value hierarchy system was static. The mechanism through which the value hierarchies are modified via messages, decay back to base values, and change those base values over time is a new addition to our agent mind that will allow us to simulate a much richer set of behaviors. We anticipate, for example, that we will be able to provide an implicit memory system for our agents by modifying base values in the hierarchy in response to events in the simulation. In an extended game comprised of a series of scenarios, this same mechanism could allow for global shifts in values in response to the actions of agents controlled by the player. Neither of these capabilities would be possible without the system we developed to support the music PMF. We will examine these and other related issues in future work.

Another result of this work is the addition of two significant PMFs to our archive. Our noise PMF is quite portable and easy to implement in any simulation that factors arousal or stress into decision-making. The music PMF is both less robust and less portable. Although it works well within our agent architecture, our implementation relies heavily on our emotion unit and could not be readily ported to other simulations that handle emotion and its effects on decision-making according to different algorithms.

The most significant result of our work, however, is the speed at which it was accomplished. Conception, research of the literature, model construction, and implementation within the architecture took less than a month from start to finish, with one programmer working full-time on the project and another half-time. This suggests that our methodology will result in relatively low development costs and excellent scalability as PMFs are added to the simulation.

7. Future Directions

The work presented here is still in an early state of development. Our simulation environment is currently little more than a demo, our agent architecture lacks the user interfaces that will allow users other than the programmers in our lab to modify agent parameters and create new agents and environments, and we have much work to do in order to optimize our agents and improve their performance.

In addition to these basic ongoing implementation issues there are a number of improvements and expansions that we are currently exploring. Perhaps the most critical of these is an effort to validate and verify our agents and the HBM/PMFs that they employ. Although we can say with a relatively high degree of certainty that our HBM/PMFs are internally valid models, we loose that certainty the moment we assemble a variety of interacting HBM/PMFs together in the same simulation. Some sort of validation is a necessity before we can move forward.

We do not expect that simulations based on our agents will necessarily be reliably predictive of real-world events, but we would like to be able to accurately simulate events that occurred in the past. For example, an analyst using our software should be able to recreate a specific event or scenario and then modify the parameters of the simulation to see how that scenario would have played itself out with different starting conditions or sequences of events. We intend to carry out an extensive array of correlational validity studies and correspondence tests to ensure that our agents' behavior closely matches that of their real-world counterparts. Based on the results of these studies we will tune and benchmark our PMFs so that other developers will have a sense of how to use them effectively.

One of our goals for our agents is that they be simulation agnostic. Ultimately, we would like simulation developers to be able to populate their own simulations with our agents. Our agent architecture has been designed with this in mind from the beginning, but to realize this goal we will need to write a sophisticated translation layer that sits in between our agents and foreign simulations and passes messages back and forth between the two. We intend to write this translation layer and then port our agent to another simulation to assess the feasibility of this approach and to determine the resources required to attempt such a port.

References

[1] Silverman, B.G., Might, R.., et al., "Toward A Human Behavior Models Anthology for Synthetic Agent Development", *10th Conf. On Computer Generated Forces and Behavioral Representation*, SISO, May. 2001.

[2] Silverman, B. G., "More Realistic Human Behavior Models for Agents in Virtual Worlds: Emotion, Stress, and Value Ontologies," Philadelphia: U of Penn/ACASA Technical Report, 2001.

[3] Silverman, B. G., Johns, M., et. al.: "Rapidly Constructing Virtual Asymmetric Opponents from Data and Models in the Literature: Case of Crowd Rioting." *11th Conf. On Computer Generated Forces and Behavioral Representation*, SISO, May. 2002. (in press).

[4] Silverman, B. G., Johns, M., et.al.: Human Behavior Models for Game-Theoretic Agents: Case of Crowd Tipping. *10th Conf. On Computer Generated Forces and Behavioral Representation*, SISO, May. 2001.

[5] Broadbent, D.E.: "Human performance and noise." In C.M. Harris (Ed.), *Handbook of Noise Control (2nd ed.)*. New York: McGraw-Hill Book Company, 1979.

[6] Donnerstein, E., & Wilson, D. W.: "Effects of Noise and Perceived Control on Ongoing and Subsequent Aggressive Behavior." *Journal of Personality and Social Psychology*. 34(5). 774-781. 1976.

[7] Edsell, R. D.: "Anxiety as a function of environmental noise and social interaction." *Journal of Psychology*. Vol 92(2), 219-226. 1976.

[8] Kruglanski, A. W., & Webster, D. M.: "Group members' reactions to opinion deviates and conformists at varying degrees of proximity to decision deadline and of environmental noise." *Journal of personality and social psychology*, 61:2. 1991.

[9] Cassity, M. D. "The influence of a music therapy activity upon peer acceptance, group cohesiveness, and interpersonal relationships of adult psychiatric patients." *Journal of Music Therapy*. Vol 13(2), 66-76. 1976.

[10] Anshel, A. & Kipper, D. A: "The influence of group singing on trust and cooperation." *Journal of Music Therapy*. Vol 25(3), 145-155. 1988.

[11] Galizio, M. & Hendrick, C.: "Effect of musical accompaniment on attitude: the guitar as a prop for persuasion." *Journal of Applied Social Psychology*, 2, 350-9. 1972.

[12] Ortony, A., Clore, G. L., & Collins, A.: The Cognitive Structure of Emotions, Cambridge University Press, Cambridge, 1988.

[13] Johns, M. & Silverman, B. G.: "How Emotion and Personality Effect the Utility of Alternative Decisions: A Terrorist Target Selection Case Study," *10th Conf. On Computer Generated Forces and Behavioral Representation*, SISO, May. 2001.

[14] Janis, I. L. & Mann, L.: *Decision Making: A Psychological Analysis of Conflict, Choice and Commitment*, Free Press (MacMillan) 1977.