

The BASP Agent-Based Modeling Framework: Applications, Scenarios and Lessons Learned

David S. Dixon & William N. Reynolds

Least Squares Software

Albuquerque, NM

ddixon@leastsquares.com

grendel@leastsquares.com

Abstract

The Behavior Action Simulation Platform (BASP) has been in existence since early 2000, when it was first applied to small-team reconnaissance scenarios for the United States Marine Corps. Since that time, BASP-based projects have seen research into new application areas, new modeling technologies and case study scenarios. The authors present some of these developments, including example models and their outcomes, and some of the lessons learned.

1. Introduction

The Behavior Action Simulation Platform (BASP) provides a general framework for agent-based modeling (ABM). Key to the BASP architecture is modularity, whereby agents can be given new capabilities by making new program modules available to the framework. These program modules, called “aspects”, need only implement a specific Java interface, and are otherwise free to perform any kind of computation or manipulation involving any accessible component of the model. Aspects that have been implemented so far include decision aspects (fuzzy logic, game theory, neural networks), action aspects (motion, detection, fire weapon), attribute aspects (position, health, weapon type), and control aspects (timekeepers, spawners, reapers, terminators).

ABM simulation is a rich field for research into complex systems where nonlinear properties emerge from the interactions themselves and it is difficult or impossible to separate the constituents. Other modeling and simulation (M&S) technologies are very strong in certain areas, for example:

Game theory – a compact representation of decisions where all the deciding factors and, more importantly, their relative values, are known in advance.

Neural Networks – can represent extremely complex decision processes as long as there is sufficient historical or generated data with which to train the network.

System dynamics – describes dynamical systems in which the components are well known and the rates of change (flow and feedback) can be expressed mathematically, such as population dynamics, resource utilization and macroeconomics.

Genetic algorithms – for cases where it is anticipated that the model will evolve (adaptation, mutation, hybridization) from one or more preceding models.

With BASP we are incorporating these technologies in ways that play to their strengths within a broader agent-based model. For example, game theory is applied to negotiations between two agents where the analyst specifies in advance what each agent brings to the bargaining table. A neural network may represent the complexities of a single powerful leader or a governing body. The network is trained either with historical data or with data from a large number of interactions where the decision process has been specified explicitly with rules, game theory or both. The trained network, then, can be exposed to new circumstances and make decisions consistent with those under which it was trained. A network is also able to learn from new experiences. System dynamics and genetic algorithm capabilities are presently under consideration.

In the following sections we will present some of our modeling experiences during the development and enhancement of BASP-based platforms.

2. The First Fuzzy Models: *Bond Market Contagion and Ethnic Conflict*

The first objective in 1999 was to develop and deploy, in a few months, an ABM of human behavior specific to a given problem but with as much reusability as possible. Fuzzy logic was chosen to represent decisions in a very human-like way, where true-false statements may also include maybe, maybe true and maybe false. The imprecision of fuzzy logic resonated with early modelers, but the fundamental process of modeling was made cumbersome by the interface. Two early adopters were able, however, to develop the following simple

models using the initial crude implementation.

Bond Market Contagion: When the Russian bond market collapsed in August of 1998, the markets of Brazil and Ecuador were hard hit even though they had no real connection to Russia. Conventional wisdom said that those two South American economies were impacted by "a reassessment of risk in emerging markets" [1] resulting from economic crises in Asia and Russia. But what did that mean in terms of specific instruments, markets, and classes of investor? Investors in emerging economies should be, after all, among the most risk tolerant. At what point were the yields no longer sufficient on high-risk instruments? Why were some vulnerable markets crushed and others were not?

The bond market model was made up of about ten investor agents, three to five bond markets, and a global economic indicator. The investor agents were given random values of assets, risk tolerance and propensity to diversify. Each market was given risk and yield values that were coupled loosely to each other, to investor confidence in that market and to the global economy. The economy fluctuated periodically and the period was adjustable. An example of one of the rules is shown in Figure 1. Exploration of the independent variables revealed the following regimes:

Steady state investor confidence – investments fluctuated between markets based on yields that were tied to fluctuations in the global economy.

Partial investor withdrawal – investors with lower risk tolerance quit the markets entirely

Abandonment of high-risk markets – all investors withdrew from the riskiest markets, moving assets to safer though lower-yield markets.

Complete abandonment – all investors withdrew from all markets.

```

Rule Name: X Bond C5   Connection Type: Investor and Country
Description: risk averse buy if rate good
Condition:
    |----your Bond Liquidity is very liquid
    |---and
    |----your Bond Rate is T-Bill + 3% to T-Bill + 4%
---and
    |----my Risk Tolerance is risk averse

Action(s):
Set Business to Category: buy
    
```

Figure 1: A Rule for the Bond Market Contagion Model. The conditional tree results in a truth-value that is used in setting the category of Business (the buy signal). A high truth-value (a real number between 0 and 1) is a strong pull to "buy".

These states appeared to be achievable for some combinations of all variables. Then the model was modified to allow the investors to communicate with each other. This amounted to a simple risk averaging, so that each individual's perception of a market's risk was related to the average of all investor's perception of risk. This addition increased the difficulty of maintaining the steady-state condition, with all

investors acting bullish in good times and bearish in the bad. This is shown in Figure 2, where movement away from buying (up) coincides with economic downturns. The abandonment of specific, high-risk markets was pronounced for rapid fluctuations of the global economy. For slower rates of change in the perceived risk of a market, the lower risk-tolerant investors pulled out early and yields were increased enough to keep the higher risk-tolerant investors in the market.

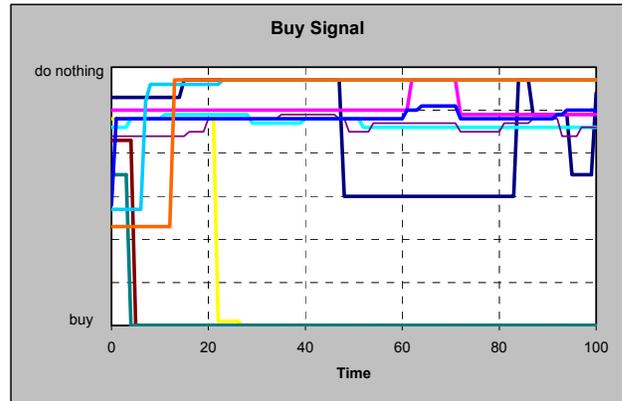


Figure 2: Buy Signal for Each of Ten Investors Over Time. The economy has a downturn every 24 cycles, each slump has random depth and lasts about 9 cycles, the first one beginning at cycle 9.

Ethnic Conflict: This scenario was an attempt to duplicate an agent-based model reported in a 1999 paper by Bhavnani and Backer [2]. The original is an elegant model with a few thousand agents, each acting based on four parameters:

Message – prevailing belief as to whether interethnic violence is localized or widespread

Trust – the extent to which the disenfranchised ethnic group trusts the ethnic group in power

Genocidal Norm – the tendency of a member of one ethnic group to justify genocide against the other

Noise – error in the Message regarding the localized or global nature of interethnic violence

From these parameters, members of the ruling group chose to attack or not attack, and members of the other group chose to run away or stand and fight. In the original, events were global and sequential: everyone got the Message, everyone interacted, and then everyone did the bookkeeping (low body count meant localized conflict, high body count meant widespread conflict) which fed back into the Message for the next round.

The effort to duplicate the original study was not completely successful, in part because of limitations imposed by purely fuzzy logic. The interpolation of fuzzy logic tended to average away the spikes in body count that would have turned the conflict from local to widespread. In an effort to

counter this, assessment of the level of violence was moved from the global context to individual agents. This had the effect of making the timing completely asynchronous, with each agent acting, assessing and feeding back into its own Message. Thus, extremists were insulated from alarmists, so that Messages from the more excitable agents weren't as infectious as they might otherwise have been.

3. Position and Movement in Fuzzy Coordinates: *George and Lenny*

It became clear from these early projects that a successful general-purpose platform would have to be organized in a way that motivated the basic questions of "*who does what to whom and why?*" without being burdened by arcane notation, esoteric jargon and a non-intuitive division into abstract objects. It was from this that the BASP design principles of accessibility, flexibility and scalability [3,4] arose in early 2000.

In the spring of 2000, the USMC's Project Albert [5] adopted BASP for a project that came to be known as Archimedes. Marine Corps analysts were interested in simulations for small team reconnaissance, focusing on behavioral factors, primarily training and discipline. Such a model would require new features: terrain, position, movement, detection, and weapons.

In the process of solving a general and "fuzzy" representation of terrain, position on the terrain, and movement to affect that position, we developed the *George and Lenny* toy model. George was a capable leader, while Lenny was a devoted follower. George was cognizant of the world around him, had goals, and some insights into how to achieve those goals. Lenny had only one behavior, to follow George, who was responsible for Lenny's wellbeing, so he waited for Lenny to catch up when necessary. Later versions of *George and Lenny* included an antagonist called Evil Hicks. George had a strong aversion to Evil Hicks, whereas Lenny, who just followed George, came into contact with Evil Hicks under some circumstances. In some scenarios, Lenny would follow the nearer of Evil Hicks and George, making it possible for Evil Hicks to capture Lenny.

The main purpose of these models was to experiment with fuzzy notions of position and movement, and to gain experience and insights into writing fuzzy logic rules for them. The fuzzy approach eliminated the need to do floating-point math to N decimal places, freeing the modeler to concentrate on general behaviors, but fuzzy logic introduced new artifacts of sensitivity ("twitchiness") and null points (indecision). One contributor to twitchiness was the familiar "heat bug" problem of agent-based modeling. For example, when four mutually interacting Lenny agents were close to George, they would buzz around him like insects. This was addressed with the realization that, when catching up to someone on the sidewalk, for example, we generally slow to

match their speed as we approach them. That meant that movement speed required feedback from distance. Another twitchiness was an artifact of fuzzy logic. A fuzzy relation (less than, greater than, etc.) returns a truth-value, a continuous variable ranging from zero (false) to one (true). The fuzzy logic rule testing if an agent has arrived at a position has a truth-value that becomes non-zero well before the agent arrives at that position. If there are other rules affecting agent behavior based on distance from the position, the two rules may arrive at their threshold truth-values in an unanticipated order, possibly resulting in oscillatory behavior. This was often solved by defining "latch" variables – two-state variables that, when tested, result in truth-values of only zero or one. In compound if-then statements, these latch variables provided truth-value inertia.

Indecision resulted from an agent being given multiple (usually two) conflicting directives that effectively cancelled. For example, in scenarios where Evil Hicks was able to capture Lenny, Lenny was sometimes caught halfway between George and Evil Hicks. George's repulsion to Evil Hicks was exactly cancelled by his responsibility to wait for Lenny. Lenny's devotion to George was exactly cancelled by his attraction to Evil Hicks. All agents became stationary or made small oscillations about a fixed point. As another example, if there were two Evil Hicks agents in George's path, his repulsion to each caused him to take a path between them. If, however, the Evil Hicks' were close together, the total repulsion was enough to bring George to a standstill. The pull of his goal was equal and opposite to the combined push of the two Evil Hicks agents. Indecision was almost always eliminated by avoiding symmetry (multiple Lenny agents with mutual interactions, for example) or with the introduction of "boredom" timers that increase continuously. For example, a timer can increase the pull of the goal over time, so that George would eventually abandon Lenny to Evil Hicks.

4. Multi-Agent Interactions: *Turkeys*

Early indications were that the behaviors of interest in an agent-based model were not in the agents themselves, but in their interactions, which in BASP are called "connections". Eventually, however, it also became clear that much of behavior – social behavior, in particular – is a matter of informing one connection from another. For example, if a motorist has a bad interaction (connection) with an auto mechanic, the motorist propagates that information in a conversation (connection) with a friend, and so on. The simplest form of multiple-connection interaction is one involving three agents (as in the example above: motorist, auto mechanic, friend). In BASP, these interactions are called "triples".

The first toy model using triples was called *Turkeys*, in which 20 to 30 agents (the "turkeys") blundered into the fire

of a machine gun emplacement. The turkeys had two directives: (a) go toward the goal, and (b) if shot, go away from the shooter. The “run away” strategy meant that every turkey was shot, some of them fatally, but the rest escaped with non-fatal injuries. (At this point, the states of shooting, being shot, and fatal versus non-fatal injuries were all variables manipulated by fuzzy rules. Later, these would become aspects.) The memory of being shot faded over time, however, so the turkeys once again attempted to reach their goal, which was on the other side of the machine gun, with predictable results. Additional toy models become progressively more sophisticated:

Repulsion – Once they discovered the machine gun, the turkeys pushed on that connection, so that repulsion to the machine gun and attraction to the goal combine to cause the turkeys to circle around the machine gun emplacement. Because of the initial injuries and fatalities, however, and subsequent fire from the flanks and rear, few of them survived to reach the goal. This is shown in Figure 3.

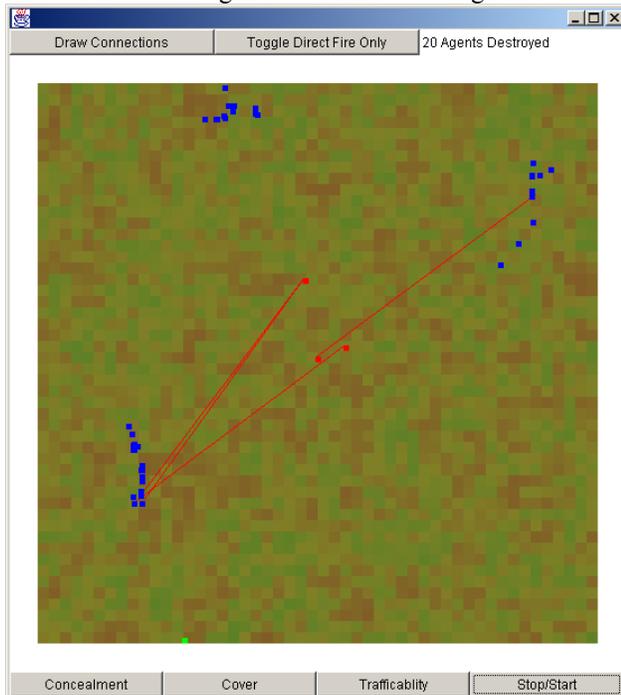


Figure 3: Turkeys - The blue dots are turkeys, the three central red dots are machine guns, the red lines are machine gun fire, and the green dot at the bottom edge is the turkeys' goal. The terrain (squares of varying shades of brown and green) is randomly generated.

Clustering – The turkeys were all interconnected so that, when the fire began and the front line of turkeys retreated they pushed the flock back with them. More turkeys survived by following the flock, but individuals remained unaware of

the machine gun until actually taking fire. The center of the flock was not repulsed from the machine gun until a majority of turkeys had received fire, but more turkeys survived than in the preceding model.

Triples – On discovering the machine gun by taking fire, turkeys on the front line spawned triples between all the other turkeys and the machine gun. This is the BASP equivalent of broadcasting information about the location of the machine gun. This way, even the turkeys that had not taken fire were now aware of the machine gun and avoided it. Most of the turkeys survived to reach the goal.

Formation – The turkeys were given small weapons that, individually, are ineffective against the machine gun. With coordination via their triples, they fell into formation, made a frontal attack on the machine gun and, though a few of them perished, their combined fire eliminated the machine gun.

The *George and Lenny* and *Turkeys* toy models served as both test benches and learning laboratories for the basic capabilities that would be required for a general-purpose combat simulation tool.

5. Fuzzy Logic and Combat Simulation: *Archimedes*

The remaining combat-specific requirements for *Archimedes* were all met in the form of aspects:

Terrain aspects – interact with the terrain to regulate speed, detection and cover, and to facilitate path-finding algorithms to maximize a weighted combination of speed, detection, and cover.

Weapon aspects – range, effectiveness, lethality (includes non-lethal weapons like sticks and rocks)

Weapon fire aspects – the act of firing a weapon and of taking fire from a weapon

Detection aspects – ability to detect enemy troops, equipment, etc. (affected by target speed, size, etc.)

Spawner and reaper aspects – create and destroy agents, connections and triples

State aspects – health, discipline, stance to support rules like "receiving weapons fire degrades health," "impaired health degrades discipline," and "lowered stance reduces detection".

The combat aspects were developed with the help of a team of Marines with both field experience and M&S expertise. The Marine analysts developed a reconnaissance scenario and rules for behavior based on Marine combat doctrine. The display for this scenario is shown in Figure 4.

In this scenario, a small team of Marines (recon) was inserted behind enemy lines in a wooded, mountainous area. A few candidate enemy missile sites had been selected from satellite intelligence, all of which the Marines must visit within a specified period of time. During this time, an enemy team (counter-recon) patrolled the region. The mission was a success if the recon team found the actual missile site and

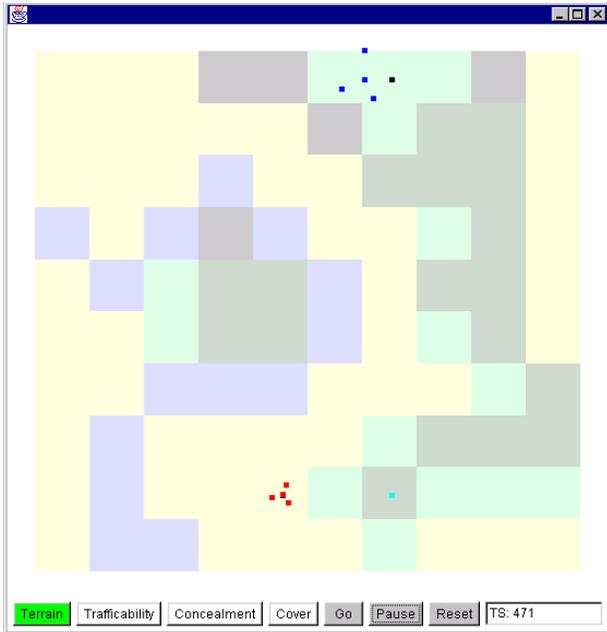


Figure 4: Display for the Reconnaissance Scenarios. Recon team is the four blue dots at the top near the black rally point (waypoint), counter-recon team is the four red dots near bottom center, and the missile site is the cyan dot to the right of bottom center. Yellow terrain is grassland, blue is river, light green is forest, black is mountain, and dark green is forested mountain.

returned to the extraction site without being detected. The mission was a failure if the recon team was detected or if it failed to find the missile site within the time allotted. For both the recon and counter-recon teams, effectiveness was a function of discipline, which included training.

During the Project Albert Summer 2000 Workshop at the Maui High Performance Computing Center in Hawai'i, this scenario was run thousands of times on parallel processors. Discipline for each team was varied across the full range, the results shown in the two-dimension plot of outcomes as a function of both teams' discipline in Figure 6. A few things to note about the plot:

- (a) As anticipated, for missions in which the recon team had higher discipline than the counter-recon team, the outcome favored the recon team and vice versa (with the following exceptions).
- (b) For a narrow band in the high discipline range for the recon team, most missions failed. This is understood to reflect excessive caution, which resulted in time running out before the mission was completed.
- (c) For very low recon team discipline, the recon team was never successful, but seldom detected, as shown by wide gaps in the X's at low recon discipline above predominantly

red triangles on the axis. This is thought to be the "dumb luck" regime, where the recon troops had insufficient training and discipline to carry out their mission at all, thereby reducing their risk of detection. The single successful outcome in the upper left quadrant is probably an example of extreme dumb luck.

(d) For low recon team discipline, scenarios resolved to failure in a consistent time of about 7000 cycles. Times to successful outcomes are quite varied until very high recon discipline, where they clustered around about 9500 cycles.

A group at West Point Military Academy developed an Archimedes model for peacekeeping in an urban terrain, while the Marine Corps undertook the rural peacekeeping model shown in Figure 5. The goal in both cases was to investigate various strategies for peacekeeping interdiction and intervention in terms of minimizing risk to peacekeeping troops, minimizing civilian casualties, and containment or capture of aggressors.



Figure 5. An Archimedes Peacekeeping Scenario. The red dots are groups of guerillas in the forest (dark green). Green dots are groups of villagers going about their business within the village. Connections between the guerillas and the villagers (red and blue lines) indicate that the parties are aware of one another. Light green and yellow represent open terrain, the sea is blue. Magenta represents a steep mountainous region.

In the example shown, when the guerillas emerged from the forest, the villagers scattered into the surrounding open terrain but some were trapped at the seashore. The guerillas then occupied the village, with occasional forays against nearby villagers. Questions asked included will the arrival of foreign peacekeeping troops before the attack prevent an incursion by the guerillas or provoke them? Would the presence of peacekeepers draw the villagers out of harm's way, or make them easier targets for the aggressors?

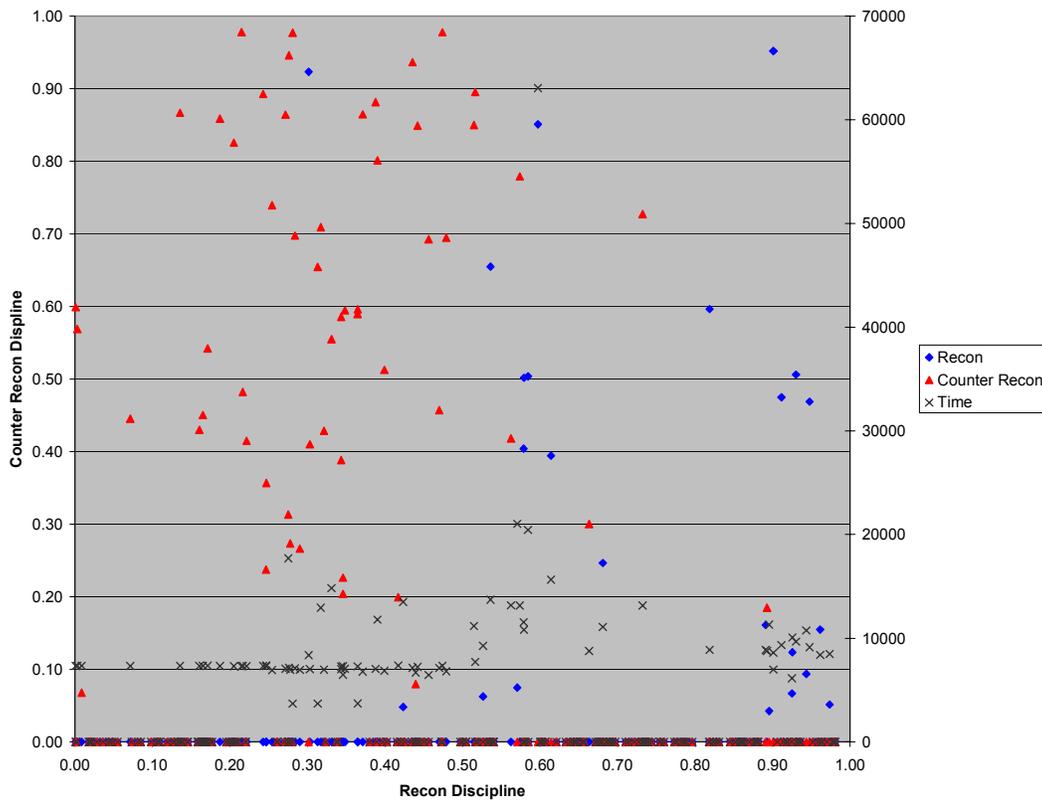


Figure 6: Approximately 1000 Archimedes Runs. Runs span the ranges of discipline for each team, counter recon along the left axis, recon along the bottom axis. Blue diamonds indicate runs in which the recon team won, red triangles indicate runs in which counter recon team detected the recon team. Shapes along the bottom axis represent scenarios in which the time limit expired (recon fails), the red shapes representing those cases in which counter recon suspected the existence but had not yet detected recon at time limit. X indicates the time, plotted against the right axis (vs. recon discipline), at which the scenario terminated (except those exceeding the time limit). The time limit was 100,000 cycles. All terminating scenarios completed within 70,000 cycles.

6. Recent Developments

Current projects will reach completion in October 2002. These projects consider socio-political models and center on negotiations and on decision-making. Example scenarios follow.

In Figure 7, there are two social, political or ethnic groups **A** and **B** and resources **R**. Each group controls some endogenous resources and there is a single shared resource. Endogenous resources may include food, strategic alliances, foreign exchange, peace, or military capabilities, for example. Examples of shared resources are oil and water. **A** and **B** negotiate and exchange resources, including access to the shared resource. The negotiations are a complex interplay of the relative

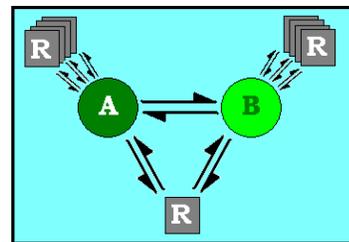


Figure 7

value of each resource as perceived by the buyer versus the value perceived by the seller. Perceived values fluctuate due to: cyclical supply, demand, and economic values of the resources; changes in the domestic

political values of resources; changes in the international political value of the resources.

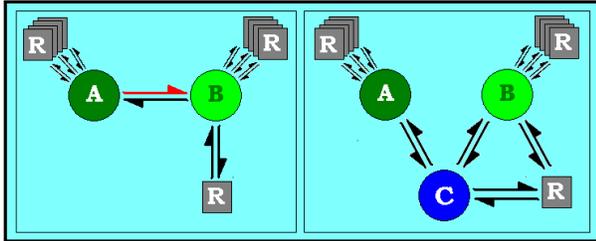


Figure 8

In Figure 8, **B** has control of a shared resource, (e.g. oil in Biafra, seaports in Croatia). In negotiations with **B**, **A** brought to the table a bid that **B** can't match: military might (red arrow). **B** responds with a bid that **A** can't match, **C**, a powerful ally. **C** may be another group like **A** and **B**, another country, or an international organization such as NATO or the UN. In this case, **C** participates actively in the political process, mediating negotiations between **A** and **B**, including **A**'s access to the shared resource. Note that, while **C** may place a high value on peace, **B**, in particular, perceives peace to be much less valuable because violence is what got **C** involved.

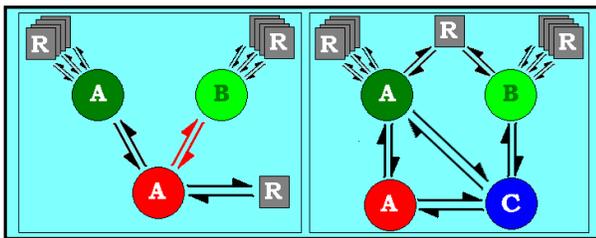


Figure 9

Figure 9 shows an alternative to the situation in Figure 8. In this case, the civilians of group **A** (upper, green circle) are supported by a sympathetic military (lower, red circle). Examples of **A** are ethnic Serbs in Kosovo or Javanese in Irian Jaya. All parties expose their respective resources and perceived values to the negotiation aspects, which mediate the political process. In this case, perhaps due to an agreement between **C** and the **A** military, control of the disputed resource is returned to joint **A** and **B** civilian control.

In Figure 10, circumstances (draught, flood, refugee influx) have made both **A** and **B** reliant on humanitarian aid, **HA**. Group **A** has attempted hostile control of **HA**, prompting the intervention of **C**. In this case, **A** does not cease hostilities, so that **C** maintains two modes of interaction with **A**, one as neutral distributor of aid, the other as military guarantor for the aid givers.

The negotiation aspect for these scenarios is being implemented using game theory. Each negotiation

(game) occurs between two players where the matrix of strategies and payoffs is based on rules defined for those

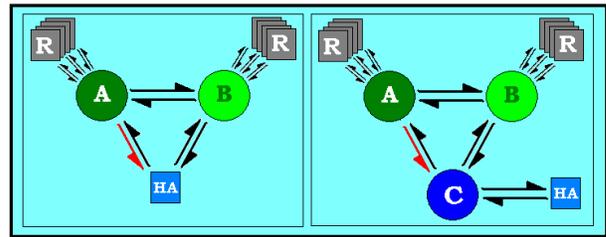


Figure 10

players in that particular interaction. Each time a negotiation is conducted, other rules may have affected the state variables from which the payoffs are computed, making it a new game from the game theoretical point of view.

7. Conclusion – Lessons Learned

One of the most interesting and consistent observations we've made throughout these various projects is that the earliest – and sometimes most significant – insights occur while reducing a problem to its most fundamental players, interactions, and basic rules of behavior. This is not to say that completing the computerized model and running simulations are wasted activities. Indeed, the insights gained in the initial phase typically direct the questions posed in simulation. This may be a corollary to astronomer Fred Hoyle's observation "...we cannot conceive of a problem until we are close to its solution." [6] In other words, often the appropriate questions to ask of a model are not known until, and are motivated by, the very act of constructing the model.

Additionally, there appears to be a law of diminishing return with regard to model fidelity. If the most important insights happen very early on, it seems also to be true that the frequency and importance of additional insights diminishes exponentially as a model is made increasingly more complex. This is an observation based on intuition and a small sample set, and very fertile ground for long-term graduate research studies some day. So far, however, we've yet to see a project where even the most basic models have been explored to their limits. For the time being, anyway, simple "toy models" provide more information and potential insights than researches have bandwidth to exploit.

Finally, the incorporation of game theory and neural networks into an agent-based framework has shown promising results. These projects, scheduled to complete at the end of 2002, show that external technologies can, in fact, be "dropped in" to BASP as aspects. This makes even more tantalizing the prospect of interoperability between BASP-based platforms and other modeling and simulation systems.

8. References

[1] C. Adams, D. J. Mathieson, G. Schinasi, "International Capital Markets Developments, Prospects, and Key Policy Issues" (World Economic And Financial Surveys, International Monetary Fund, September 1999), p. 63,
<http://www.imf.org/external/pubs/ft/icm/1999/pdf/file03.pdf>

[2] R. Bhavnani, D. Backer, *Journal of Conflict Resolution*, (Department of Political Science, Yale University, June 2000),
<http://www.yale.edu/unsy/jcr2jun2000.htm>

[3] W. N. Reynolds, D. S. Dixon, paper presented at Complex Systems and Policy Analysis: New Tools for a New Millennium,

RAND Science & Technology Policy Institute, Arlington VA, Sept 27-28, 2000, <http://www.leastsquares.com/papers/rand2000.pdf>

[4] W. N. Reynolds, D. S. Dixon, in *Maneuver Warfare Science 2001*, G. Home, M. Leonardi, Eds. (USMC, Quantico, VA, 2001),
<http://www.leastsquares.com/papers/mws2001.pdf>

[5]
http://www.mors.org/meetings/C4ISR_2000/Tutorials/Project_Albert.pdf

[6] F. Hoyle, *Astronomy and Cosmology* (Freeman, San Francisco, 1975), p. 685.