

Distinguishing Control and Plant Dynamics in Enterprise Modeling

H. Van Dyke Parunak¹, Sven Brueckner¹, John Sauter¹, and Robert S. Matthews¹
ERIM Center for Electronic Commerce, PO Box 134001, Ann Arbor, MI 48113-4001 USA
{vparunak, sbrueckner, jsauter, rmatthews}@erim.org

Abstract

Two factors can confound the interpretation of an enterprise model. First, the dynamics of the control technology interact in complex ways with those of the plant, and engineers need to be able to distinguish these effects. Second, “mean field” approximations of the behavior of the system may be useful for qualitative examination of the dynamics, but can differ in surprising ways from the behavior that emerges from the interactions of discrete agents. This paper examines these effects in the context of a specific research project applying agent-based control to a military air operations scenario.

1. Introduction

Modern military operations can overwhelm a commander. The information available from satellite and other sensors floods conventional analysis methods. Enemy forces using advanced technology can hide or change location faster than conventional planning cycles can respond, and coordinating central orders across thousands of friendly resources can slow response even further. These features are prototypical of many modern enterprise control problems.

The ADAPTIV project (Adaptive control of Distributed Agents through Pheromone Techniques and Interactive Visualization) applies fine-grained agent techniques to the control of air resources charged with defending a friendly region from enemy attack. Intelligence on the location and strength of enemy (“Red”) resources leads to the deposit of synthetic pheromones [1, 5, 7] in a spatial model of the battlespace. The propagation and evaporation of these pheromones model the uncertainty in the available intelligence. Friendly (“Blue”) units then move in response to the flow field generated by these pheromones.

Initial experiments with these mechanisms are being conducted in a scenario dealing with the suppression of enemy air defenses, a task known as

SEAD. Red is moving ground troops (GT), under cover of air defense units (AD), toward blue’s territory. Blue has bombers (BMB) to hold back the ground troops, but these bombers are vulnerable to red AD. Blue also has fighters (SEAD) that can suppress the AD and thus protect BMB. The scenario takes the form of a strategy game, called “SEADy Storm,” played on a hexagonally tiled field. Units of both sides can decide how to move over this field, and (when they find themselves in the same hexagon as one or more enemy units) whether to engage the adversary. Outcome rules determine how the strength of each unit changes as the result of an engagement.

Our initial experiments with pheromone mechanisms in this environment yield surprisingly complex variations in outcome as we vary the distribution of resources on each side. Some of this variation results from the simple strategies being executed by our agents, but much is due simply to the complexity of the game rules. To separate the two effects, we have done three successive layers of abstraction away from the initial experimental setting, first neutralizing the effect of Blue strategies, next removing the spatial structure of the game entirely, and finally abstracting away from the individual units with a mean field approximation to the combat rules.

We conducted these experiments to explore the potential of pheromone mechanisms, but this paper does not discuss these mechanisms. Our focus here is on two more general morals. First, experimental abstraction is a useful (even necessary) technique for understanding the relative effects of agents and environment (or to use the vocabulary of control theory, of the controller and the plant). Second, modeling technology can distort the picture in ways to which the analyst must be sensitive.

Section 2 describes in more detail the problem domain, the particular scenario we are exploring, and the initial results we obtained. Section 3 looks at three successive abstractions we applied to this scenario in an effort to distinguish environmental and agent

dynamics. Section 4 discusses our experience and summarizes key insights

2. The SEADy Storm Experimental Context

First we summarize the structure and rules of our plant. Then we describe the behavior of our control agents, and show results from initial experiments.

2.1 The SEADy Storm Game

SEADy Storm [3] is a war game used to explore technologies for controlling air tasking orders. The battlespace is a hexagonal grid of sectors, each 50 km across (Figure 1). Friendly (Blue) forces defend a region in the lower left against invading Red forces that occupy most of the field. Red’s playing pieces include ground troops (GT’s) that are trying to invade the Blue territory, and air defense units (AD’s, surface-to-air missile launchers) that protect the GT’s from Blue attack. Blue has bombers (BMB’s) that try to stop the GT’s before they reach the blue territory, and fighters tasked with suppressing enemy air defenses (SEAD’s).

Each class of unit has a set of commands from which it periodically chooses. In our implementation, ground-based units (GT and AD) choose a new command once every 12 hours, while air units (BMB and SEAD) choose once every five minutes. These times reflect the time it would take the resource to move across a sector. The commands fall into three categories (Table 1). GT cannot attack Blue forces, but can damage BMB’s if they attack GT.

Blue can attack AD and GT when they are moving

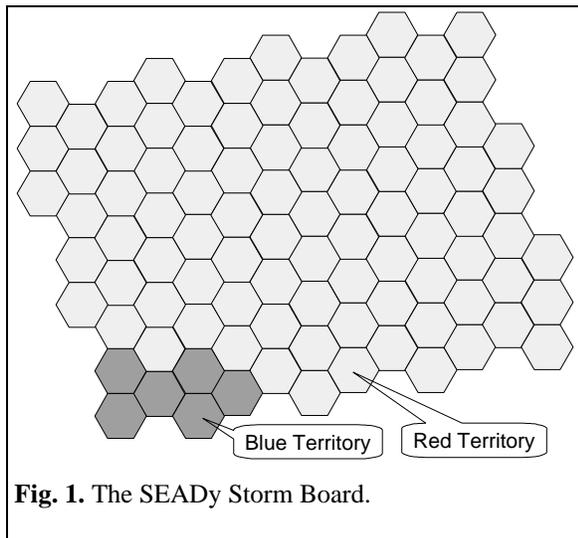


Fig. 1. The SEADy Storm Board.

Table 1. Unit Commands in SEADy Storm.

	Move	Attack	Wait
AD	Relocate	Fire (on any Blue aircraft)	Hide Deceive
GT	Advance		Hide
SEAD	NewSectors	AttackAD	Rest
BMB	NewSectors	AttackAD AttackGT	Rest

or attacking, and AD may attack any Blue forces that are not moving or waiting. Each unit has a strength that is reduced by combat. The strength of the battling units, together with nine outcome rules, determine the outcome of such engagements. Informally, the first five rules are:

1. Fatigue: The farther Blue flies, the weaker it gets.
2. Deception: Blue strength decreases for each AD in the same sector that is hiding.
3. Maintenance: Blue strength decreases if units do not rest on a regular basis.
4. Surprise: The effectiveness of an AD attack doubles the first shift after the unit does something other than attack.
5. Cover: BMB losses are greater if the BMB is not accompanied by enough SEAD.

Rules 6-9 specify the percentage losses in strength for the units engaged in a battle, on the basis of the command they are currently executing. For example, Rule 9, in full detail, states: “If BMB does “AttackGT” and GT does “Advance”: a GT unit loses 10% for each BMB unit per shift; a BMB unit loses 2% per GT unit per shift.”

2.2 The ADAPTIV Mechanisms

ADAPTIV uses pheromone techniques to control Blue operations in such a scenario. Intelligence reports on Red locations deposit synthetic pheromones in a spatial model corresponding to the hexagonal grid. The pheromone infrastructure *evaporates* pheromones to model the decreased value of stale information, *propagates* pheromones from one sector to another to pass information to nearby resources, and *aggregates* deposits from subsequent reports to highlight new information. The experiments reported in this paper manipulate a package of BMB and SEAD as a unit.

Each class of unit must wait (5 minutes for a BMB-SEAD package, 12 hours for AD and GT) after executing one command before choosing another. When a unit is eligible for a new command, it selects with equal probability from its possible commands. If it selects a movement command, the movement depends on the unit in question and the pheromones it senses in the six adjacent sectors. A unit “follows” the

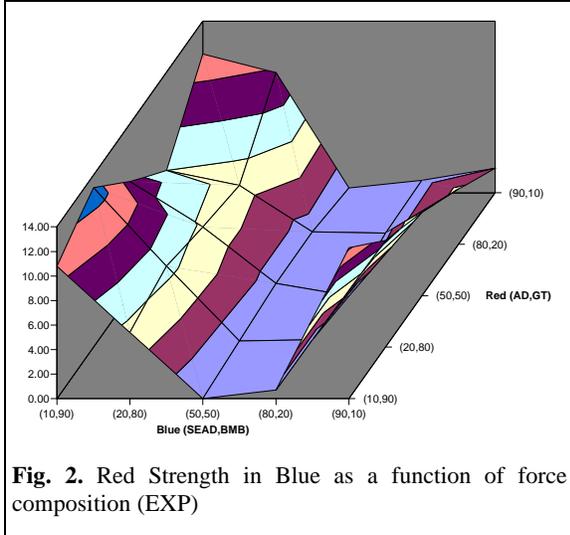


Fig. 2. Red Strength in Blue as a function of force composition (EXP)

pheromone field using a roulette wheel weighted by the strength of the appropriate pheromone in the neighboring sectors. In this experiment,

- the SEAD-BMB package follows GT pheromones,
- AD units follow a product of BMB and GT pheromones, thus seeking out BMB's that are threatening GT's, and
- GT units move randomly, with higher weights given to south-westerly neighbors.

An experiment runs for 1200 simulated hours (about 7 minutes on a 500 MHz Wintel). At the end of the experiment we calculate the total Red strength that has reached Blue territory ("Red in Blue" or "RinB") and the surviving percentages of each class of unit. We run each configuration of parameters eleven times with different random seeds, and report medians for each configuration.

2.3 EXP: Experimental Results

The primary parameter explored in the experiments reported here is the proportion of SEAD in the Blue military, and of AD in the Red military. Each side began with a 100 units, each with unit strength, and 10%, 20%, 50%, 80%, or 90% of

SEAD or AD. The uneven spacing reflects a basic statistical intuition that interesting behaviors tend to be concentrated toward the extremes of percentage-based parameters. In current military doctrine, 50% is an upper limit on both AD and SEAD. We explore higher values simply to characterize the behavioral space of our mechanisms.)

The central outcome is total Red strength in Blue territory at the end of the run (Figure 2). The landscape shows several interesting features, including

- a "valley" of Blue dominance for all Red ratios when Blue SEAD is between 50% and 80%, with slightly increasing Red success as the AD proportion increases;
- clear Red dominance for lower SEAD/BMB ratios, decreasing as SEAD increases;
- a surprising increase in Red success for the high SEAD and low AD levels.

Figure 3 shows the percentages of each class of unit surviving at the end of the run. The AD, GT, and BMB plots reflect the main features of the topology. The increase in Red effectiveness for high SEAD appears to be due to a drop in BMB survival in this region, a surprising effect since BMB's have strong SEAD protection here.

The overall system shows interesting and non-trivial dynamics, with two sources: the pheromone-based movement of the resources, and the outcome rules that define the scenario. Early study, for example, showed us that the Red superiority for low

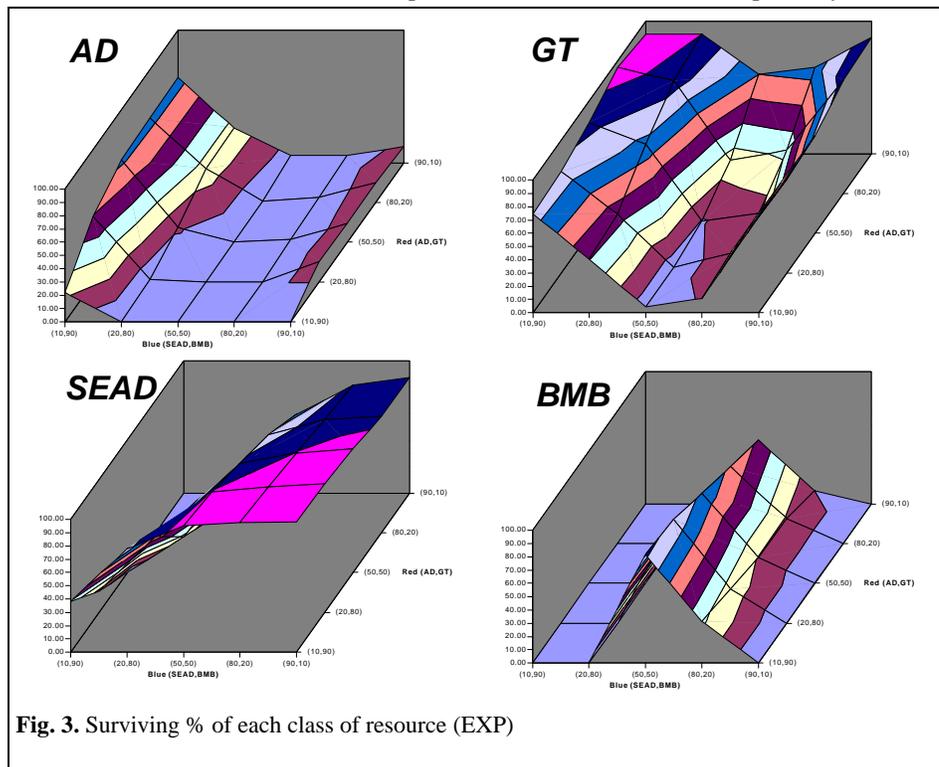


Fig. 3. Surviving % of each class of resource (EXP)

SEAD ratios is directly related to Rule 5, which places a particularly heavy penalty on Blue packages that do not have at least one SEAD for every two BMB's. This rule induces a threshold nonlinearity at SEAD/BMB 33/67, which marks the edge of the Blue valley in other runs (not shown here) that explore the parameter space in more detail. However good Blue's pheromone algorithms are at finding and targeting Red troops, Rule 5 will impose a performance cliff along this parameter.

3 Successive Abstractions

To understand the contribution of our control mechanisms, we must distinguish their dynamics from those of the plant with which they interact. We abstract away successive details of our mechanisms and compare the resulting system behaviors with those of the full system. First, we performed a mean field abstraction (ABS3) that removes the effects of Blue strategy, spatial distribution, and the distinction among individual agents. This abstraction behaved sufficiently differently from EXP that we examined two intermediate abstractions, one removing only blue strategy (ABS1), the other removing blue strategy and spatial distribution (ABS2). We present the abstractions in logical sequence rather than in chronological order.

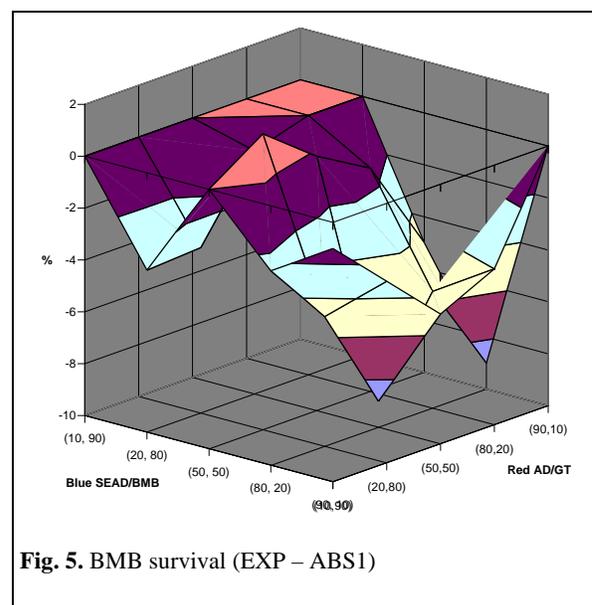
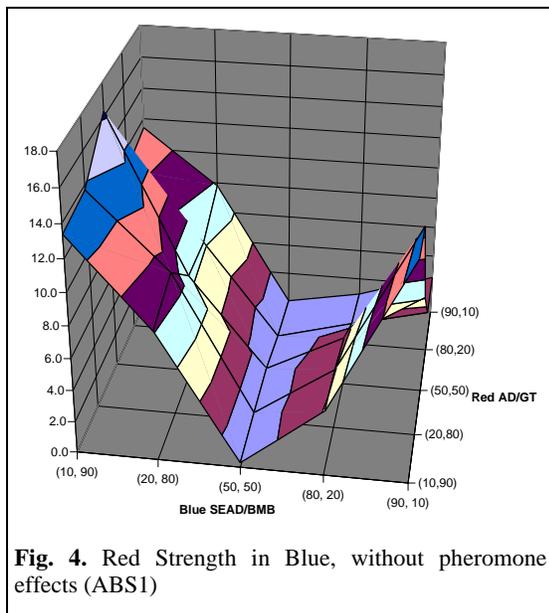
3.1 ABS1: Ignoring Blue Strategy

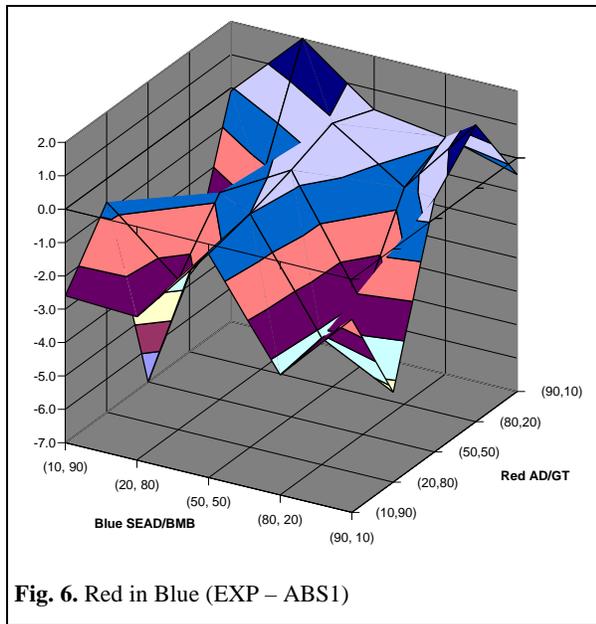
Blue units find Red targets by climbing pheromone gradients. A logical abstraction is to "cut off their

noses," moving Blue randomly rather than in response to pheromone signals. Figure 4 shows the strength of Red in Blue at the end of the game under these conditions. The landscape has the same general features as Figure 2.

We can compare the two by subtracting at each point the Red in Blue strength when Blue moves randomly from that when Blue follows pheromones, as in Figure 5. Because Blue seeks to keep Red out of Blue territory, differences less than 0 represent a net contribution of the Blue mechanisms. Figure 5 shows that our mechanisms are generally effective, with the greatest benefit at 10% SEAD, 50% AD. There are two exceptional regions where random wandering outperforms pheromones.

The first is when Red AD is above 50%. In this region, BMB survival is worse in EXP than in ABS1 (Figure 6), leading us to hypothesize two possible causes for the difference. It may result from the movement rule we have assigned to AD, to follow GT weighted by BMB pheromones. If BMB movement is regular (guided by slow-moving GT), AD's can position themselves more effectively. When BMB move randomly, AD's have a harder time positioning themselves for maximum impact. Another explanation recognizes that with high AD coverage around GT, a frontal attack of BMB on GT takes them into the most deadly opposition, while random movement will sometimes encounter weaker Red units that they can more effectively pick off. Distinguishing these alternative effects requires further experiments.





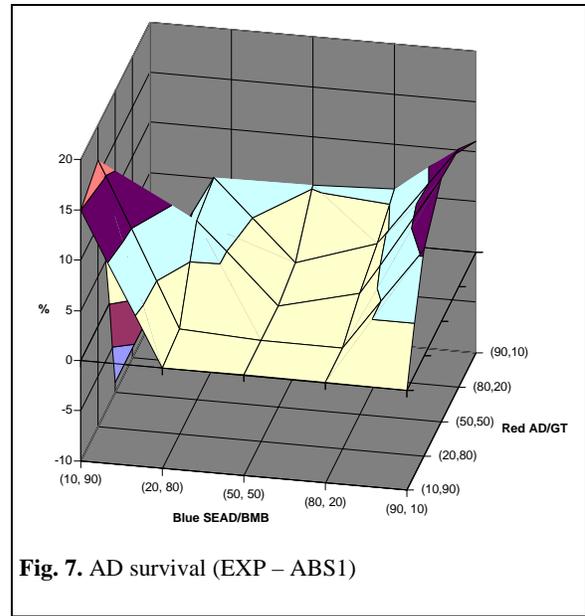
Lesson: Within the same problem domain, parametric differences can lead to a very different interaction between the agents and their environment, and agents need to be designed to take these differences into account in determining their behavior.

Second, even with low Red AD, the pheromone mechanisms make little or no contribution at 50% SEAD. This effect is probably due to a lack of opportunity. The valley around 50% SEAD and low AD is intrinsically so favorable to Blue that Blue's strategy makes little difference.

Lesson: Success may result from luck rather than intelligence. Environmental dynamics can be so strong that agent intelligence makes little or no difference.

Figure 7 shows the difference in AD survival between EXP and ABS1. As we have seen, pheromones make little difference toward the low-AD end of the 50% SEAD valley, and help Blue at 10% SEAD, 50% AD. But they are a detriment around the edges of the valley. The right-hand ridge reflects the puzzle we have already seen in Figure 2: why should higher SEAD strength lead to better Red success? Figure 7 shows that this anomaly is reflected in AD survival. Initially, this circumstance is even more puzzling than high Red in Blue in this region. Why should higher SEAD help AD survival? By focusing our attention on the AD units, this plot leads us to the answer, a complex chain of interlocking events.

1. Rule 9 specifies that when BMB attacks GT, the losses on each side depend on the ratio of GT to BMB. When SEAD percentage is high, a package contains fewer BMB's, and the GT/BMB ratio is higher. For very high SEAD levels, BMB is at a disadvantage in an encounter with GT, accounting for the lower survival of BMB at high SEAD.



2. AD's are attracted by GT pheromones, weighted by BMB pheromones. The BMB population falls off at high SEAD levels, both intrinsically and because of the dynamic in the previous point. Thus AD movement becomes more random, and AD's are less likely to be found in the close vicinity of GT.
3. SEAD's travel with BMB's in Blue packages, which are attracted by GT. As AD's wander more, they are less likely to be near GT's, thus less likely to encounter SEAD's, and their survival increases.
4. Meanwhile (returning to the right-hand flap in Figure 2), the decreased population of BMB's leaves GT free to invade Blue territory.

Lesson: Even simple rules interact in complex and unanticipated ways. Careful analysis is necessary to understand the implications of these interactions for the design of individual agent behaviors.

3.2 ABS2: Ignoring Spatial Distribution

ABS1 shows us the effect of turning off Blue's pheromone mechanisms, but Red's deliberate movement is another layer that we must peel away from the system behavior to understand the impact of the outcome rules. One way to remove the effect of these mechanisms would be to randomize Red's movement as well as Blue's, but this would still leave a dependency on Red's initial spatial distribution. Alternatively, we can remove space entirely, so that all units occupy a single sector.

Historically, we analyzed ABS2 in order to validate the mean field approach of ABS3, which is intrinsically non-spatial, so ABS2 pursues the single-sector approach. This abstraction changes how Red

and Blue agents encounter one another. In the spatial model, agents interact when they find themselves in a common sector. As a result, agent movement (whether random or purposeful) induces a distribution on how many agents can be engaged at a given time step. For example, Blue in a sector with no Red forces can neither cause nor receive battle damage. Under such circumstances, an “attack” command is effectively a no-op. When we place all resources in the same sector, we need another way to model how many resources will be engaged. Thus we define the proportion of each type of unit that will execute each eligible command at each time step. In the results reported here, we assign the following parameters (parameter set a), based on the results from this same set in ABS3:

- AD: 0% Hide, 0% Deceive, 10% Fire, 90% Relocate (to the same sector)
- GT: 80% Hide, 20% Advance (to the same sector)
- SEAD: 10% Rest, 90% Attack AD, 0% New Sector
- BMB: 60% Rest, 20% Attack AD, 20% Attack GT, 0% New Sector

For example, at a given time step, a randomly selected 80% of the GT’s will Hide, while the others will Advance (thus being vulnerable to attack).

Figure 8 summarizes some results from these parameters, compared with EXP and ABS1. These plots show several interesting features.

The topography in ABS2-a is shifted toward lower SEAD percentages, relative to that in EXP and ABS1. The valley in surviving GT and the peak in surviving BMB now fall between 20% and 50% SEAD, rather than at or beyond 50% SEAD as before. The location of the valley reflects the penalty imposed by Rule 5 when the ratio of SEAD to BMB falls below 1/2. In ABS1, SEAD and BMB are packaged based on the overall percentage of SEAD, which is thus involved in any combat. In ABS2-a, the proportion of SEAD and BMB in a conflict depends not only on the overall percentage of SEAD, but also on the number of each that is resting and out of action on a given cycle. The command percentages in Figure 8 make 90% of SEAD available to attack on any given cycle, but only 40% of BMB. Thus the effective SEAD percentage is more than twice the overall SEAD percentage, and ABS2-a shows the same effect at 20% SEAD that EXP and ABS1 show at 50% SEAD.

The obvious fix is to fit the command percentages more carefully to the distribution induced by movement in the spatially distributed case. Such a fit is more easily requested than delivered. The

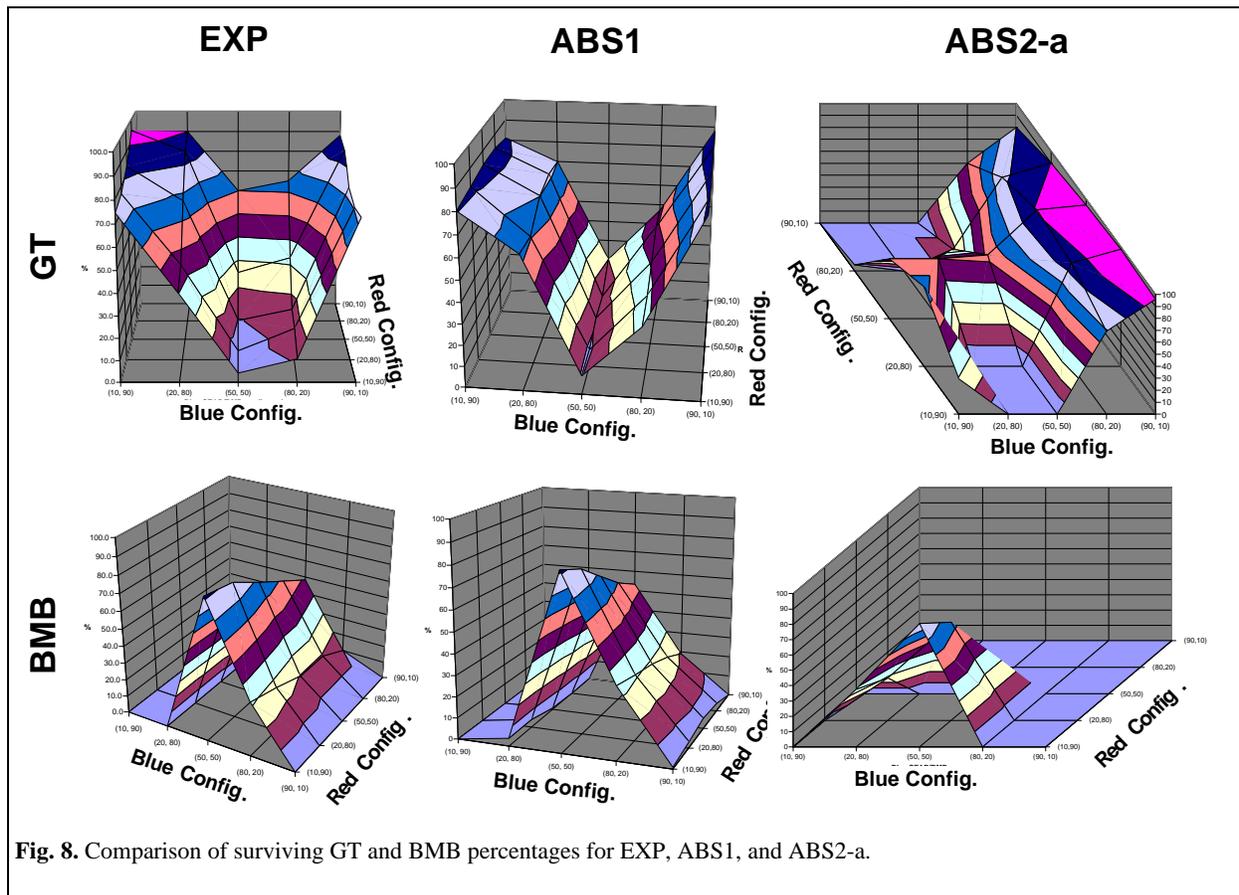


Fig. 8. Comparison of surviving GT and BMB percentages for EXP, ABS1, and ABS2-a.

complexity of various movement rules makes an analytical derivation intractable. The desired distribution is probably not even stationary, since population changes over a run will change the probability that two units will encounter one another. One could use a visual or statistical match between performance landscapes such as those in Figure 8 to determine command percentages experimentally. More fundamentally, these observations call into question the validity of aspatial models of spatially distributed problems.

Lesson: Space is not just a neutral medium in which agents interact. It plays an active and complex role in their interactions, a role that is difficult if not impossible to capture without modeling space directly.

ABS2-a also differs from EXP and ABS1 in the nonlinearity of GT's dependence on AD percentage. In the previous experiments, the valley rises monotonically as AD increases. In ABS2-a, this rise peaks at 50% AD, then falls sharply for higher AD percentages. At this point, we do not have a detailed explanation for this feature. It is unlikely that we will devote considerable effort to understanding it, since the basic lesson of ABS2-a is that removing spatial distribution entirely from the model is not a fruitful approach to our objective of factoring agent effects from environment effects.

3.3 ABS3: Ignoring Unit Identity

Execution of agent-based models can be very time-consuming (in our case, requiring 7 minutes to simulate 1200 hours). Eleven replications at each of $5 \times 5 = 25$ Red/Blue configurations require over 32 hours to yield experimental results. A parallel equation-based model (ABS3) requires about 1.5 seconds to simulate 1200 hours. Significant differences between agent-based and equation-based models [8, 10] make the agent-based model the gold standard for evaluating our pheromone methods, but rapid surveys of parameter space with an equation-based model might guide more meticulous (and time-consuming) verification using the agent-based model.

ABS3 uses a population-based modeling approach where the aggregated strength of all units of one type (AD, GT, SEAD, BMB) is represented in the size of one distinct population. The size of a population changes over time. The change is determined by the portion of each population that engages in combat in each discrete time step and by the losses inflicted in these combats.

We represent the population dynamics in a set of difference equations that capture both, the combat composition and the outcome rules. For example, in

the GT population the population size is reduced by in every step by

$$\Delta GT = g(GT, Advance, BMB, AttackGT) \\ * BMB * \frac{c(t, BMB, AttackGT)}{c(t, GT, Advance)}$$

where

- $g(X, a, Y, b)$ represents the percent losses of a group of units of type X that executes the command a when it engages a group of units of type Y that executes the command b at a time step. Losses are specified in outcome rules five to nine.
 - $c(t, X, a)$ specifies the combat composition, and represents the percentage of population X that executes the command a at time t . The ABS3 experiments all assume a constant combat composition: $c(t1, X, a) = c(t2, X, a)$ for all pairs $(t1, t2)$.
- ABS3 initializes the four populations to represent the initial strength of the combatants. Then, for a specified number of time steps, it
- computes the decrease in the population size for each population;
 - limits the computed losses to the portion of each population that is actually engaged in combat (given by: $X * c(t, X, a)$, where a is an attack command); and
 - applies the losses.

The number of time steps is the same as the number of calls in a comparable run of EXP to the function resolving combat situations.

With ABS3, we were able to generate landscapes that matched those from EXP and ABS1 qualitatively, but not in detail. The differences might be explained either by the move from an agent model to an equation one, or by the collapse of space. We constructed ABS2 in an effort to tease apart those rival effects.

As we have seen, collapsing space does make a difference, due largely to the necessity to capture in static command probabilities the distribution of activities induced by agent encounters as they move through space. Comparison of ABS2 with ABS3 shows that the move from agents to equations has other effects as well.

Figure 9 compares three pairs of surviving GT and BMB landscapes. ABS3-a uses parameter set a (defined in Section 3.2), and shows that even with a space-free model, command percentages can be tuned to produce landscapes similar to those in ABS1 (or, for that matter, EXP; compare Figure 3). We ran ABS2-a with these same percentages to test whether the shift to an equation-based model makes a difference. Figure 9 shows that it does. The percentages that produce realistic landscapes in ABS3-a lead to the anomalies we have already

discussed in ABS2-a. The third column in Figure 9 shows landscapes in ABS3-b, with different command parameters chosen to make these landscapes resemble ABS2. Parameter set b is $AD = \{0.0,0.0,0.5,0.5\}$; $GT = \{0.8,0.2\}$; $SEAD = \{0.45,0.55,0.0\}$; $BMB = \{0.2,0.4,0.4,0.0\}$.

Thus ABS3 can show us the existence of interesting non-trivial performance landscapes, but for a given set of parameters, it cannot reliably tell us either the location or the topology of their features. The salient difference between ABS2 and ABS3 is that ABS2 retains distinct agents, while ABS3 represents only the aggregate strength of the entire population of agents of a given type (thus, a single strength for each of AD, GT, SEAD, and BMB). The strength of individual agents in ABS2 evolves from the engagements in which each agent is involved, and thus summarizes that agent's history. ABS3 loses this history. The simulation logs show different evolution of the total strength over time in the two cases, leading to the different final outcomes reflected in Figure 9. The effect is closely related to the sensitive dependence of nonlinear systems on initial conditions. Once individual agents in ABS2 come to differ slightly in their strengths, their subsequent evolution can diverge greatly, leading to changes in the outcome of subsequent combats. ABS3 cannot track these different histories, and so is insensitive to their results.

Lesson: Like spatial distribution, ontological distribution (distributing processing over independent interacting processes) makes a substantive and non-intuitive difference to the outcome of a model. Whether or not a mean-field model works in a given situation is an empirical question. Such models must be carefully validated against agent-based models before being trusted.

4 Discussion and Summary

The world is too complicated a place to understand as it exists. Science seeks to understand it by abstracting away details to leave a simplified system, and manipulating that system with a modeling technology (such as mathematical analysis or simulation). The validity of this process requires that neither the abstraction nor the modeling representation substantially change the behavior of the system. As our experience shows, both of these requirements are easily compromised.

First, abstraction requires that the system under study be cleanly separable from the environment in which it is usually embedded. Designers of agent-based systems typically pay much more attention to the agents than to the environment. The complex

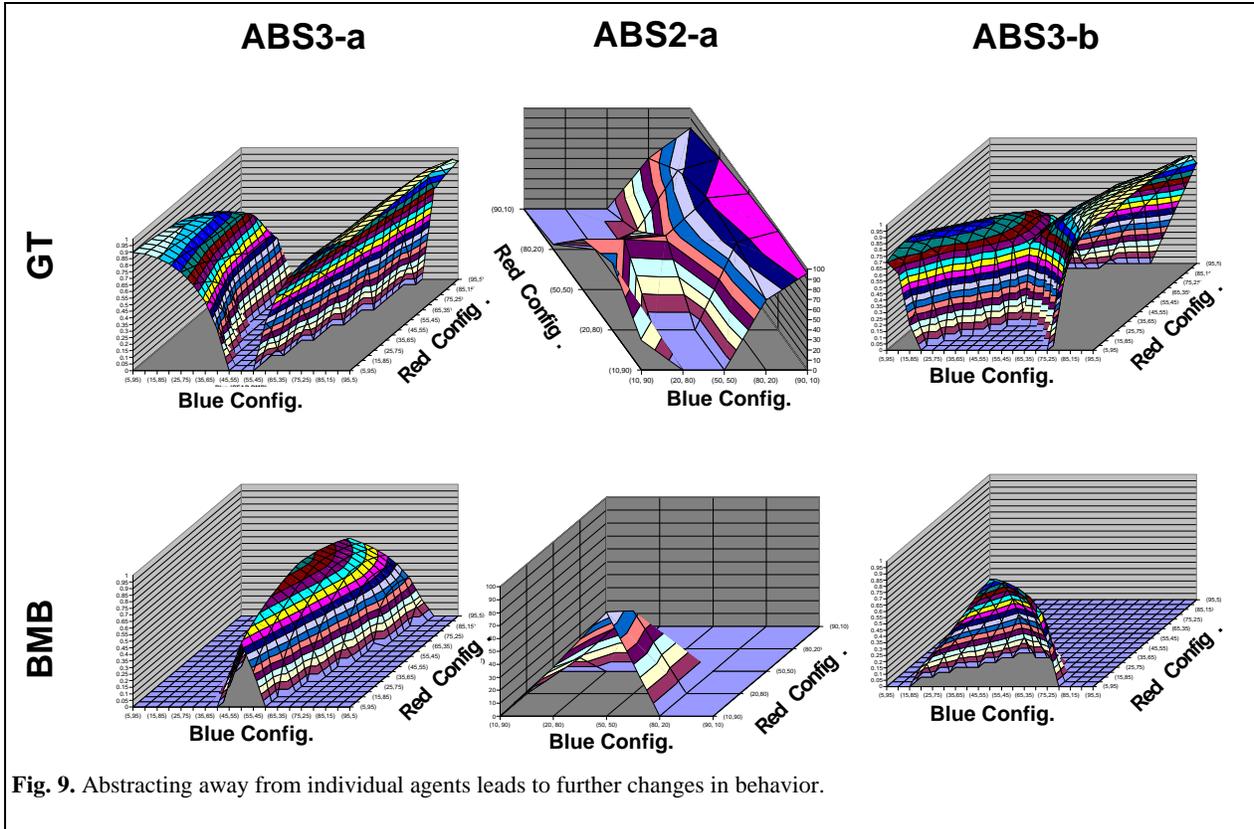


Fig. 9. Abstracting away from individual agents leads to further changes in behavior.

interactions discussed in this paper show that the environment deserves more attention. Our observations support researchers in embodied cognitive science [9] who argue that the agent and its environment must be designed together. The behavior of interest is that of the whole system, and only by considering the environment with the agent can we reliably design systems that do what we wish. The abstraction process exemplified in this paper is a methodological tool that can make us aware of how our systems interact with their environments.

Second, modeling technologies are not content-neutral. Their mechanisms can introduce artifacts determined more by the modeling technology than by the system being modeled. Earlier researchers have pointed out such effects within agent-based models, based on differences between synchronous and asynchronous execution [2, 4]. Our results in this paper reinforce our earlier observations [8] about the loss of ontological distribution in an equation-based model.

Our conclusion is cautionary, not fatalistic. We do not reject system modeling and simulation as impossible. In fact, it is unavoidable in engineering control systems for complex enterprises, due to the analytical intractability of typical systems and their strongly nonlinear behavior [6]. We do warn that simulation is nontrivial. It forms a new way of doing science, alongside physical experimentation and mathematical analysis. These classical modes have evolved methodological guidelines for reliable results. Effective simulation science requires the development of similar guidelines, and the particular potential of agent-based modeling suggests that agent researchers should be in the forefront of developing this methodology. This paper is a step in this direction.

Acknowledgements

Olga Gilmore and Murali Nandula of the ERIM CEC staff contributed significantly toward the experimentation reported in this paper. This work is supported in part by the DARPA JFACC program under contract F30602-99-C-0202 to ERIM CEC. The

views and conclusions in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the US Government.

References

- [1] S. Brueckner. *Return from the Ant: Synthetic Ecosystems for Manufacturing Control*. Ph.D. Thesis at Humboldt University Berlin, Department of Computer Science, 2000.
- [2] B. A. Huberman and N. S. Glance. Evolutionary Games and Computer Simulations. *Proceedings of the National Academy of Science USA*, 90(August):7716-7718, 1993.
- [3] A. Kott. SEADy Storm. 2000. Web Page, <http://www.cgi.com/web2/govt/seadystorm.doc>.
- [4] S. E. Page. On Incentives and Updating in Agent Based Models. *Computational Economics*, 10:67-87, 1997.
- [5] H. V. D. Parunak. 'Go to the Ant': Engineering Principles from Natural Agent Systems. *Annals of Operations Research*, 75:69-101, 1997. Available at <http://www.erim.org/~van/gotoant.pdf>.
- [6] H. V. D. Parunak. From Chaos to Commerce: Practical Issues and Research Opportunities in the Nonlinear Dynamics of Decentralized Manufacturing Systems. In *Proceedings of Second International Workshop on Intelligent Manufacturing Systems*, pages k15-k25, K.U.Leuven, 1999.
- [7] H. V. D. Parunak and S. Brueckner. Ant-Like Missionaries and Cannibals: Synthetic Pheromones for Distributed Motion Control. In *Proceedings of Fourth International Conference on Autonomous Agents (Agents 2000)*, 2000. Available at www.erim.org/~vparunak/MissCann.pdf.
- [8] H. V. D. Parunak, R. Savit, and R. L. Riolo. Agent-Based Modeling vs. Equation-Based Modeling: A Case Study and Users' Guide. In *Proceedings of Workshop on Multi-agent systems and Agent-based Simulation (MABS'98)*, pages 10-25, Springer, 1998. Available at <http://www.erim.org/~van/mabs98.pdf>.
- [9] R. Pfeifer and C. Scheier. *Understanding Intelligence*. Cambridge, MA, MIT Press, 1999.
- [10] W. G. Wilson. Resolving Discrepancies between Deterministic Population Models and Individual-Based Simulations. *American Naturalist*, 151(2):116-134, 1998.