Distinguishing Environmental and Agent Dynamics: A Case Study in Abstraction and Alternate Modeling Technologies

H. Van Dyke Parunak[1], Sven Brueckner[1], John Sauter[1], and Robert S. Matthews[1]

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 MAS application or model. First, the MAS's dynamics interact in complex ways with those of its environment, and agent engineers need to distinguish the two. Second, "mean field" approximations of the behavior of the system may be useful for qualitative examination of the dynamics, but can differ surprisingly from the behavior that emerges from the interactions of discrete agents. This paper examines these effects in the context of a project applying synthetic pheromones for 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 characteristics are hallmarks of situations for which agent-based systems are particularly suitable.

ADAPTIV (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 (SP's) [3, 12, 14] in a spatial model of the battlespace. The propagation and evaporation of SP's model the uncertainty in the available intelligence, and generate a flow field that guides friendly ("Blue") units.

In our experiments with these mechanisms, Red moves ground troops under cover of air defense toward Blue's territory. Blue's bombers can attack the ground troops, but are vulnerable to Red air defenses. Blue's fighters can suppress air defenses and thus protect bombers. Our initial experiments with SP's in this scenario 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 applied three successive abstractions to 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 develop two morals that go beyond SP's. 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 briefly motivates and reviews SP technology, explaining our mechanisms and comparing our approach with other related work. Section 3 describes the problem domain, our experimental scenario, and initial results. Section 4 applies three successive abstractions to this scenario in an effort to distinguish environmental and agent dynamics. Section 5 discusses our experience and summarizes key insights.

2. Synthetic Pheromones for Multi-Agent Coordination

Many applied problems require that entities move from one location to another under certain constraints. These problems have traditionally been addressed with centralized planning and control mechanisms. In highly dynamic domains such as combat management, central coordination may not permit timely response, and the central controller is a vulnerability that can place the entire operation at risk.

Negotiation schemes from recent MAS research enable the entities to maintain such constraints. These mechanisms effectively decentralize system control, but can require sophisticated inter-agent communication and significant processing within individual agents. There are several motives for seeking simpler mechanisms.

- The resources needed for conventional negotiation restrict the deployment of these techniques to highly numerous, relatively inexpensive entities such as boxes of field rations or seat assemblies en route from supplier to final assembly.
- The dynamics of classical DAI mechanisms are poorly understood, and in some cases may become intractable, leading to inadequate performance.
- The complexity of designing interlocking protocols so that they all work correctly is often non-trivial.
- In a military context, high semantic content in inter-agent messages is a point of vulnerability to evesdropping.

Insect colonies perform sophisticated motion coordination and control using neither central coordination nor direct agent-to-agent communication. We are developing mechanisms that mimic their behavior in engineered systems.

2.1. Insect Examples

Insects perform impressive feats of coordination without direct inter-agent coordination, by depositing pheromones (chemical scent markers) in the environment and then sensing them [12]. For example, ants construct networks of paths that connect their nests with available food sources. Mathematically, these networks form minimum spanning trees [8], minimizing the energy ants expend in bringing food into the nest. Graph theory offers algorithms for computing minimum spanning trees, but ants do not use conventional algorithms. Instead, this globally optimal structure emerges as individual ants wander, preferentially following food pheromones and dropping nest pheromones if they are not holding food, and following next pheromones while dropping food pheromones if they are holding food.

Brownian motion brings the ant arbitrarily close to every point in the plane. As long as the separation between nest and food is small compared with the ant's range, a wandering ant will find food if there is any, and a food-carrying ant will find the nest.

Because only food-carrying ants drop food pheromone, and because ants carry food only after picking it up at a source, all food pheromone paths lead to food. Because only empty ants drop nest pheromone, and because all empty ants originate at the nest, all nest pheromone paths lead home. Because pheromones evaporate, paths to depleted food sources disappear, as do paths laid down by ants that get lost.

The initial path will not be straight, but the tendency of ants to wander even in the presence of pheromones will generate short-cuts across meanders. Overlapping pheromone paths tend to merge together into a trace that becomes straighter the more it is used. The character of the resulting network as a minimal spanning tree is not intuitively obvious from the individual behaviors, but emerges from the emulation.

2.2. Mechanisms

These behaviors manifest two mechanisms that we seek to emulate: stochastic movement, and pheromones.

Stochastic search is the ultimate "weak method." By itself, it fails under combinatorial explosion. However, when tempered with other mechanisms (such as temperature in simulated annealing, or crossover in genetic algorithms, or local alignment with other agents in particle swarm optimization), it is ubiquitous in practical weak methods. It is an essential component in most models of insect behavior, and we will incorporate it in our model.

The real world provides three operations on chemical pheromones that support purposive insect actions. It *aggregates* deposits from individual agents (providing integration of information across multiple agents and through time), *evaporates* them over time (thus forgetting obsolete information and avoiding overloading), and *diffuses* them to nearby places (generating a gradient that agents can follow). A *pheromone infrastructure* [3] is a software environment that supports these operations. It consists of a network of places over which agents move, and between which pheromones propagate. At any moment in time, an agent is located at a specific. This place offers it a number of services, including the ability deposit pheromones and query the strength of pheromones deposited by others.

These techniques can be applied to real-world problems in two ways. Sometimes actual physical entities can move immediately in response to pheromones. If physical entities cannot tolerate the intrinsic stochasticity, the emergent behavior of a population of virtual agents can be consulted to guide the real entities. These simple mechanisms can run extremely rapidly, permitting simulation in faster than real-time.

2.3. Comparison with Other Research

The potential of insect models and similar mechanisms for multi-agent coordination and control is receiving increasing attention [1, 12]. [5, 7] offer theoretical discussions with simple applications, and [6] shows how these techniques can play a credible game of chess. [19] constructs a synthetic chemistry through which agents interact.

Pheromone-based techniques have been developed for routing telecommunications packets [2] and moving physical entities [15]. The latter work appealed to a neural backpropagation model for its antecedents. However, the accumulation of weights on frequently activated links through backpropagation has many formal similarities to the accumulation of pheromones on well-traveled paths. Steels proposed similar mechanisms for coordinating small robots used in exploring remote planets [18]. These techniques are applicable to a range of optimization problems including the traveling salesperson problem and the quadratic assignment problem [4].

Our approach is distinct in three ways.

1. We extend SP's with mechanisms that permit human overseers to monitor and influence them as they operate.
2. We hybridize SP's subsymbolic reasoning with symbolic processing.
3. We give special attention to tools and methods for engineering SP's for real-world problems. This paper's theme reflects this distinctive.

3. The SEADy Storm Experimental Context

First we summarize the experimental scenario. Then we describe the behavior of our agents, and show the results from our initial experiments.

3.1. The SEADy Storm Game

The SEADy Storm war game [10] exercises technologies for controlling air tasking orders. The battlespace is a hexagonal grid of 50-km sectors (Figure 1). Friendly (Blue) forces defend against invading Red forces that include ground troops (GT's) that are trying to invade the Blue territory, and air defense units (AD's) that protect the GT's from Blue attack. Blue has bombers (BMB's) to stop the GT's,, and fighters to suppress enemy air defenses (SEAD's).

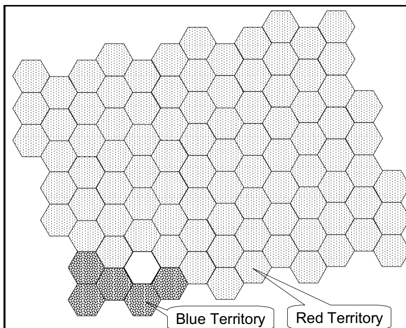

Fig. 1. The SEADy Storm Board.

Each class of unit has a set of commands from which it periodically chooses. 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,

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

Table 1. Unit Commands in SEADy Storm.

reflecting unit velocities. 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 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."

3.2. The ADAPTIV Mechanisms

ADAPTIV controls Blue operations with pheromone techniques. Intelligence reports on Red locations deposit SP's in a spatial model reflecting the hexagonal grid. The experiments reported here manipulate a package of BMB and SEAD as a unit. When a unit is eligible for a new command, it selects with equal probability from its possible commands. If it selects a movement command, its movement depends on its class and the SP's it senses in its own and the adjacent sectors. A unit "follows" the pheromone field using a roulette wheel weighted by the strength of the SP's:

- 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. Metrics include 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.

3.3. 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.

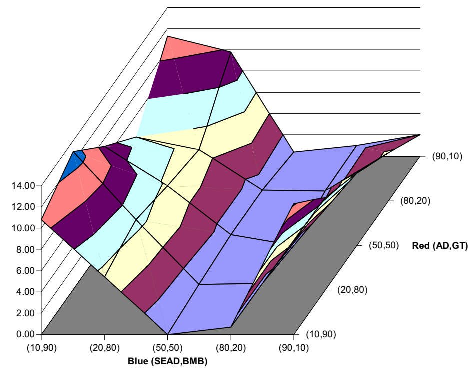

Fig. 2. Red Strength in Blue as a function of force composition (EXP)

Figure 3 shows the surviving percentages of each class of unit at the end of the run. AD, GT, and BMB 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.

These interesting and non-trivial dynamics have two sources: the pheromone-based movement of the resources, and the outcome rules that define the scenario. For example, Red superiority at low SEAD ratios is directly related to Rule 5, which

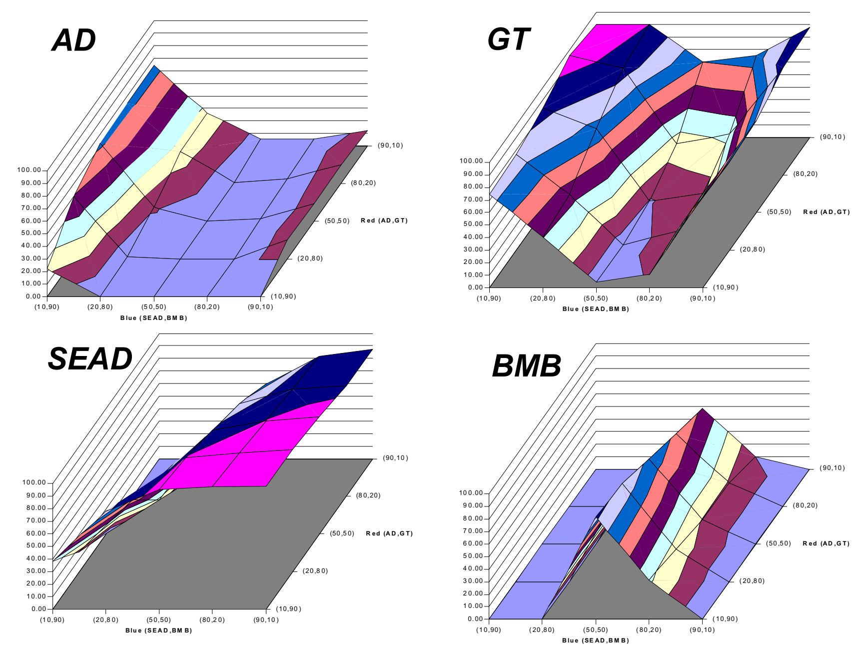

Fig. 3. Surviving % of each class of resource (EXP)

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.

4. Successive Abstractions

To understand the contribution of our mechanisms, we must distinguish their dynamics from those of their environment. 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. Because this abstraction behaved differently from EXP, 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. Due to space limitations, we discuss only those details that illustrate the impact of our successive abstractions.

4.1. 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

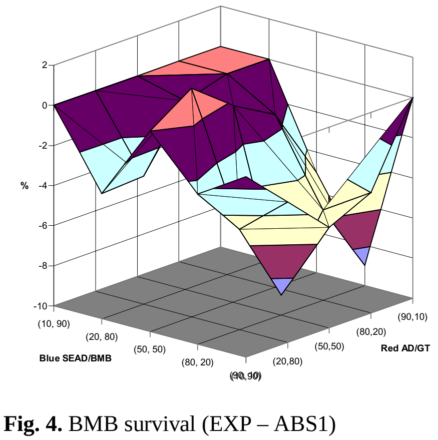

Fig. 4. BMB survival (EXP – ABS1)

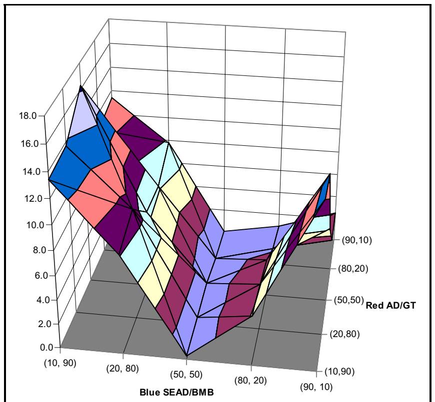

Fig. 5. Red Strength in Blue, without pheromone effects (ABS1)

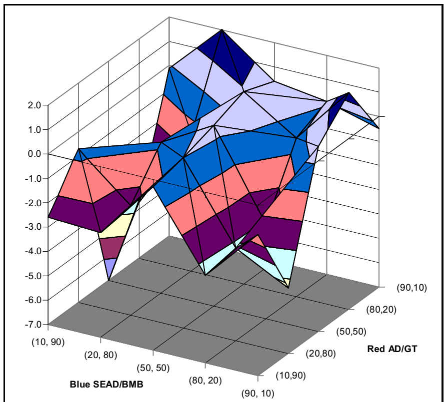

Fig. 7. Red in Blue (EXP – ABS1)

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. Designers of agents need to take these differences into account.

Second, even with low Red AD, SP mechanisms make little or no contribution at 50% SEAD, probably due to a lack of opportunity. The valley around 50% SEAD and low AD is 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,

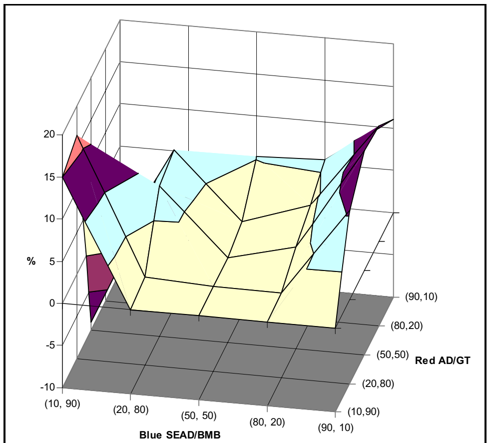

Fig. 6. AD survival (EXP – ABS1)

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 GT/BMB. With high SEAD, 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 that designers must seek to understand through careful analysis.

4.2. 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.

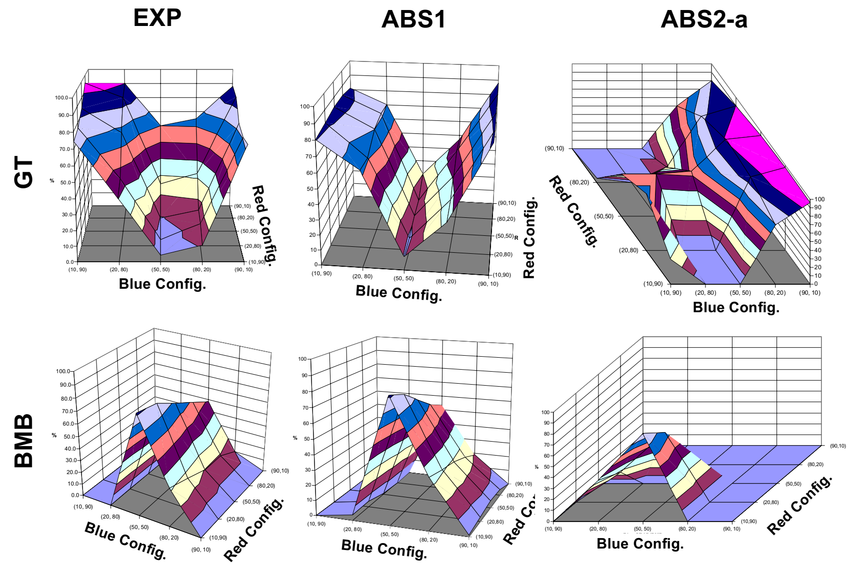

Fig. 8. Comparison of surviving GT and BMB percentages for EXP, ABS1, and ABS2-a.

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.

An 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 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.

4.3. ABS3: Ignoring Unit Identity

Execution of agent-based models can be time-consuming (in our case, 7 minutes for 1200 hours). Eleven replications at each of 5x5 = 25 Red/Blue configurations require over 32 hours. A parallel equation-based model (ABS3) requires about 1.5 seconds to simulate 1200 hours. Significant differences between agent-based and equation-based models [16, 20] 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 slower 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 theses 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.

ABS3 yields landscapes that match 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. ABS2 teases 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 GT and BMB landscapes. ABS3-a uses parameter set a (defined in Section 4.2), 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 (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

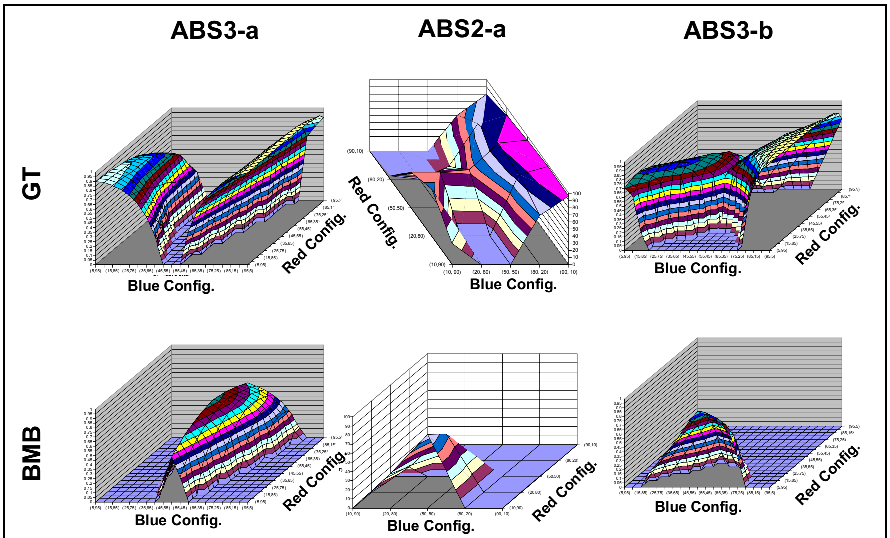

Fig. 9. Abstracting away from individual agents leads to further changes in behavior.

over independent interacting processes) substantially and non-intuitively affects a model's results. 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.

5. 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, designers of agent-based systems typically pay more attention to the agents than to the environment. The complex interactions discussed in this paper show that the environment deserves more attention. Our observations support researchers in embodied cognitive science [17] 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. They 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 [9, 11]. Our results in this paper reinforce our earlier observations [16] about the loss of ontological distribution in an equation-based model.

System modeling and simulation are not impossible. In fact, they are unavoidable in engineering agent-based systems, due to the analytical intractability of typical systems and their strongly nonlinear behavior [13]. However, 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.

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] E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. New York, Oxford University Press, 1999.

[2] E. Bonabeau, F. Henaux, S. Guérin, D. Snyers, P. Kuntz, and G. Theraulaz. Routing in Telecommunications Networks with "Smart" Ant-Like Agents. Santa Fe Institute, Santa Fe, NM, 1998.

[3] S. Brueckner. *Return from the Ant: Synthetic Ecosystems for Manufacturing Control*. Thesis at Humboldt University Berlin, Department of Computer Science, 2000.

[4] M. Dorigo, V. Maniezzo, and A. Colorni. The Ant System: Optimization by a Colony of Cooperating Agents. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 26(1):1-13, 1996.

[5] A. Drogoul. *De la Simulation Multi-Agents à la Résolution Collective de Problèmes*. Thesis at University of Paris IV, 1993.

[6] A. Drogoul. When Ants Play Chess (Or Can Strategies Emerge from Tactical Behaviors? In *Proceedings of Fifth European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW '93)*, pages 13-27, Springer, 1995.

[7] J. Ferber. *Multi-Agent Systems: An Introduction to Distributed Artificial Intelligence*. Harlow, UK, Addison Wesley Longman, 1999.

[8] S. Goss, S. Aron, J. L. Deneubourg, and J. M. Pasteels. Self-organized Shortcuts in the Argentine Ant. *Naturwissenschaften*, 76:579-581, 1989.

[9] 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.

[10] A. Kott. SEADy Storm. 2000. Web Page, http://www.cgi.com/web2/govt/seadystorm.doc.

[11] S. E. Page. On Incentives and Updating in Agent Based Models. *Computational Economics*, 10:67-87, 1997.

[12] H. V. D. Parunak. 'Go to the Ant': Engineering Principles from Natural Agent Systems. *Annals of Operations Research*, 75:69-101, 1997.

[13] 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.

[14] 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.

[15] H. V. D. Parunak, J. Kindrick, and B. Irish. Material Handling: A Conservative Domain for Neural Connectivity and Propagation. In *Proceedings of Sixth National Conference on Artificial Intelligence*, pages 307-311, American Association for Artificial Intelligence, 1987.

[16] 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.

[17] R. Pfeifer and C. Scheier. *Understanding Intelligence*. Cambridge, MA, MIT Press, 1999.

[18] L. Steels. Cooperation between Distributed Agents through Self-Organization. Vrije Universiteit Brussel AI Laboratory, 1989.

[19] T. White and B. Pagurek. Towards Multi-Swarm Problem Solving in Networks. In *Proceedings of Third International Conference on Multi-Agent Systems (ICMAS'98)*, pages 333-340, IEEE Computer Society, 1998.

[20] W. G. Wilson. Resolving Discrepancies between Deterministic Population Models and Individual-Based Simulations. *American Naturalist*, 151(2):116-134, 1998.