Forthcoming in Proceedings of ICCDM 2010, Miami, FL, July 2010

CHAPTER 1

Multi-Perspective, Multi-Future Modeling and Model Analysis

H. Van Dyke Parunak and Sven A. Brueckner
Vector Research Center, division of TechTeam Government Solutions

INTRODUCTION

Many simulation models are ad hoc, constructed to answer a specific question with a main emphasis on a limited semantic domain (such as social interactions, geospatial movement, or progress toward accomplishment of a task). This approach limits the reusability of individual models, and makes the construction of large, complex models inefficient.

This paper discusses three characteristics of our new systems approach to modeling.

Multi-future: We need to compute the envelope of possible behaviors with probability distributions over the relevant state variables, rather than generating only a single system trajectory with each run.

Multi-perspective: We must be able to integrate different model dimensions, including geospatial, social, process, financial, communications, and political, in a single model. For such a model to be computationally tractable and understandable, these different perspectives should be implementable as distinct modules that can be interconnected according to the needs of a specific model.

Analyzable: We need a set of mathematical tools that can express and test objective propositions about model behavior, instead of relying on subjective impressions formed from human observation of runs.

A modular **multi-perspective** modeling technology allows the straightforward combination and reuse of different specialized models into more complex systemic models. A **multi-future** approach guards us from unjustified generalizations based on limited numbers of runs that undersample the behavioral space of a model, and provide the basis for a statistically grounded **analytic** toolbox that can assess the significance and confidence to be associated with model outputs.

First we discuss a novel agent-based modeling formalism, the polyagent, that allows the efficient sampling of multiple futures. Then we show how this formalism supports multi-perspective modeling, and finally discuss the analytic framework such a system enables.

POLYAGENTS AND MULTI-FUTURE MODELING

THE PROBLEM

Imagine $n + 1$ entities in discrete time. At each step, each entity interacts with one of the other n. Thus at time t its interaction history $h(t)$ is a string in n^t. Its behavior is a function of $h(t)$. This toy model generalizes many domains, including predator-prey, combat, innovation, diffusion of ideas, and disease propagation.

It would be convenient if a few runs of such a system told us all we need to know, but this is not likely to be the case.

- We may have imperfect knowledge of the agents' internal states or details of the environment. If we change our assumptions about these unknown details, we can expect the agents' behaviors to change.
- The agents may behave non-deterministically, either because of noise in their perceptions, or because they use a stochastic decision algorithm.
- Even if agents' reasoning and interactions are deterministic and we have accurate knowledge of all state variables, nonlinear decision mechanisms or interactions can result in chaotic dynamics, so that tiny differences in individual state variables can lead to arbitrarily large divergences in agent behavior. A nonlinearity can be as simple as a predator's hunger threshold for eating a prey or a prey's energy threshold for mating.

An Equation-Based Model (EBM) typically deals with aggregate observables across the population. In the predator-prey example, such observables might be predator population, prey population, average predator energy level, or average prey energy level, all as functions of time. No attempt is made to model the trajectory of an individual entity.

An Agent-Based Model (ABM) describes the trajectory of each agent. A single run of the model captures only a subset of possible interactions among the agents.

In our general model, during a run of length τ, each entity will experience one of n^τ possible histories (a worst-case estimate, since domain constraints may make many of these histories inaccessible). The population of $n + 1$ entities will sample $n + 1$ of these possible histories. It is often the case that the length of a run is orders of magnitude larger than the number of modeled entities ($\tau >> n$).

Multiple runs with different random seeds is only a partial solution. Each run only samples one set of possible interactions. For large populations and scenarios that permit multiple interactions on the part of each agent, the number of runs needed to sample the possible alternative interactions thoroughly can quickly become prohibitive. In one recent application, $n \sim 50$ and $\tau \sim 10{,}000$, so the sample of the space of n^τ possible entity histories actually sampled by a single run is vanishingly small. We would need on the order of τ runs to generate a meaningful sample, and executing that many runs is out of the question. Previous multi-trajectory approaches (Gilmer and Sullivan 1998; Gilmer and Sullivan 2000) essentially replicate single trajectory runs, and the sampling approach they take to avoid explosion in the search space leads to an ergodicity failure (Gilmer and Sullivan 2001).

THE POLYAGENT MODELING CONSTRUCT

In the polyagent modeling construct (Parunak and Brueckner 2006), a persistent *avatar* manages a stream of transient *ghosts*, each of which explores an alternative future for the entity in a simulated world. These futures are executed in one or more virtual *environments*, such as a book of temporally successive geospatial maps or a task network, whose topology reflects that of the problem domain.

Ghosts are tropistic. Their behavior is determined by a set of fields in their environment (called “digital pheromones” after the insect mechanism that inspired them). Each field associates a scalar value with each cell of the environment. Some fields are emitted by objects of interest (such as roads or buildings). Others are deposited by the ghosts as they move about. A ghost’s behavior is determined by a weighted sum of the pheromones it senses in its vicinity, where the weights define the ghost’s personality and can be either manually coded or learned by observation of the entity that the ghost represents.

In the real world, one entity’s behavior can depend on the presence or absence of another entity. A ghost’s behavior depends on the fields of other entities, and thus reflects an average response across all of the locations of the other entities that their ghosts have explored (FIGURE 1.1).

Each ghost has a state, which changes to reflect its interaction with the environment. For instance, its strength may change as a result of combat in a battlefield model. The ghost’s strength can be interpreted as its degree of health, or more abstractly as the probability that the entity that it represents would be at full strength at the ghost’s time and place. A log of each ghost’s strength as a function of time is an additional resource, alongside the pheromone fields deposited by the ghost, for deriving fields. An entity’s avatar can estimate the strength of its entity by taking the average of the strengths of its ghosts.

Each ghost increments the field corresponding to the entity it represents as it moves. The strength of this particular field at a location represents how frequently ghosts of that entity visit that location. The amount of the deposit depends on the ghost’s strength, so its field takes into account the effects of attrition. A ghost can increment multiple fields. For example, one might correspond to its own avatar, one to the entire team to which its avatar belongs, and one to a unit within the team. A field modulated by strength yields an estimate of the probability of encountering a unit of force at each location. Fields can also be modulated by a ghost characteristic than strength, such as current preference for a given course of action, yielding a

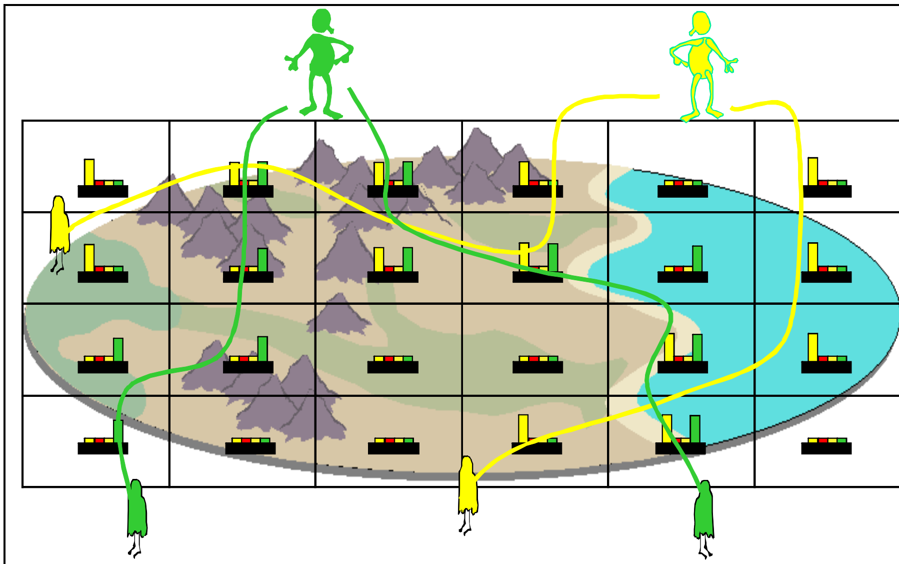

FIGURE 1.1 Ghosts explore alternate futures for their avatars, and interact via fields that they maintain in the environment.

field with different semantics that may be useful in some applications.

Field strength depends not only on entity type and location, but also on time. In predictive applications, we maintain a set of field maps, one for each successive time step from a specified time in the past (the "insertion horizon") to a specified time in the future (the "prediction horizon"). Each page covers the entire area of interest (AOI). This set of maps is called the "book of maps." The number of pages is fixed, and as real time advances, we drop the oldest page and add a new page one time step further into the future. Pages are indexed by τ. $\tau = 0$ corresponds to "now." $\tau < 0$ indexes pages in the past (used to train the ghosts by evolution against observations (Parunak 2007)), and $\tau > 0$ indexes pages in the future. Thus if the current real time is t, the real time represented on a given page with index τ is $\tau + t$.

Pages in the book of maps for which $\tau \geq 0$ have real fields only for relatively persistent environmental features such as topography or clan territories. Otherwise, the fields to which ghosts respond on these pages are built up by the ghosts themselves as they traverse them. The first ghosts to visit each page do not see any ghost-generated fields, and their behavior is constrained only by persistent features. To enable ghosts to respond to one another, avatars release them in *shifts*. In one application, each avatar releases a total of 200 ghosts over 100 shifts, two per shift. The ghosts in each shift respond to the state of the fields as modified by the previous shifts. Ghosts in early shifts do not have well-defined fields to which to respond, so their movements are not as reliable an estimator of entity movement as those in later shifts, when the fields have converged. To accommodate this increase in accuracy over time, at each simulation step the field strengths on each page are attenuated by a constant factor E (a process inspired by pheromone evaporation in insect systems). The effect is to weight deposits by later shifts more strongly than those by earlier ones. When the shifts are complete, the avatars take their action in the real world based on the information from their ghosts, we index the book of maps, and start a new avatar cycle.

Often we are interested not in the location of an individual entity, but in the distribution of a group of entities (for example, all members of a team). In this case, all ghosts of entities on the same team increment the same field, which now reflects the probability of encountering any team entity over the area of interest.

MULTIPLE FUTURES FROM POLYAGENTS

How many different futures does this approach explore? Each avatar sends $g*\sigma$ ghosts, where g is the number of ghosts issued per shift. One recent application used $n = 5$ avatars, each sending out $g = 2$ ghosts per shift over $\sigma = 100$ shifts into a book of maps containing 60 future pages. Each ghost could in principle follow a distinct path through the book of pheromones. In practice, environmental constraints mean that many ghosts follow similar paths, and our probability distributions reflect the resulting distribution of trajectories.

A state of the world consists of the state of all avatars. Because we can capture multiple avatar states concurrently, we capture a number of states of the world equal to the product of the number of states visible for each avatar. Naively, we might estimate the number of possible futures as $(g*\sigma)^n \sim 3.2*10^{11}$. This is an overestimate, for two reasons:

- The number of ghosts that have visited a given page depends on the page. Pages further in the future see fewer ghosts from the current avatar cycle. (Depending on the configuration, ghosts from earlier cycles may still be on the page.)
- A ghost interacts with later ghosts by way of the field that it increments, and this field evaporates over time. So we should not count all ghosts equally.

Assume that we are at shift σ and page $\tau < \sigma$, so that the page in question has been visited. The oldest deposit on page τ was made by the g ghosts issued at $\sigma = \tau$, and a fraction $g*E^{\sigma-\tau}$ remains. The most recent deposit, made at σ, contributes g. So each avatar's "virtual presence" on the page is

$$g\sum_{i=0}^{\sigma-\tau} E^i = g\frac{1-E^{\sigma-\tau}}{1-E} \quad (1)$$

In the near future (pages 10 and lower), we are exploring nearly 40 alternative behaviors for each avatar. The concave nature of the function is felicitous: the drop-off in parallelism is gradual until we get to more distant futures, where the prediction horizon effect (Parunak, Belding et al. 2007) (the increasingly random divergence of future trajectories under nonlinear iteration) makes predictions less reliable anyway.

The number of states explored on each page is this value raised to the power of the number of entities. Averaging this value over the 60 pages yields an average number of parallel futures of 8.4×10^7. This is several orders of magnitude lower than the 3.2×10^{11} estimate based on 200 ghosts per avatar, but still far more than a single-trajectory simulation can explore.

MULTI-PERSPECTIVE MODELS

THE PROBLEM

Modern analysis must consider multiple perspectives on the world, including geospatial, social, process, financial, communications, and political. Concurrent analysis of such a set of systems face three major challenges: combinatorics, nonlinearity, and model maintenance.

Combinatorics.—Let network i be $<N_i, E_i>$, where N is a set of nodes and E a set of edges. The set of places at which n networks can influence each other is the set product

$$\prod_{i=1}^{n} N_i \qquad (2)$$

without even taking into account the types of edges (typically more than one type per network).

Nonlinearities.—The dynamics of processes on even a single network are highly non-linear. For example, even a random graph, an oversimplified model of almost any network, exhibits a phase transition as the probability of node connection increases (Erdös and Rényi 1960). These nonlinearities mean that network predictions are subject not only to state uncertainty (resulting from ignorance about the network), but also to dynamic uncertainty (resulting from the divergence of trajectories). Nonlinear systems are notoriously resistant to closed-form analysis, and are commonly studied using simulation, but dynamic uncertainty makes single simulation runs uninformative, and the space of possible interactions is so large that even replications of conventional simulations grossly undersample it.

Model Maintenance.—Conventional modeling techniques require close integration among the various model perspectives. It is difficult to construct and test components independently, then easily reconfigure them into a variety of combinations.

HOW POLYAGENTS HELP

A given entity exists simultaneously in multiple environments. Conventional modeling focuses attention on the environment, but we focus attention on the agent, which as an Inter-Tier Entity (FIGURE 1.2) links the environments together (Parunak, Brueckner et al. 2010). The agent's experiences in one environment change its behavior in another.

To achieve a **multi-perspective** model, an agent travels from one topology to another in the course of its execution. For example, if a method in a process graph requires geospatial movement, an agent on a task node of the process graph detours through the geospatial topology to estimate the time required to complete the task. Each agent interacts with other entities by means of the fields that their agents generate, effectively interacting with all of the futures explored by those agents.

FIGURE 1.3 shows the interaction of a process model (consisting of a series of actions with ordering constraints among them) and a geospatial model. The process model specifies that the agent purchases something at its initial location, then moves to one of two alternative locations, where it seeks to sell its purchase. In a pure process simulation, the duration and success of the "Move" actions would be sampled from a distribution. In our approach, when an agent reaches a "Move" action, it drops down into the geospatial environment and actually executes the move. It may encounter obstacles or other agents in that space that affect the success or duration of its move. If it completes the move, it returns to the process model and advances to the next task.

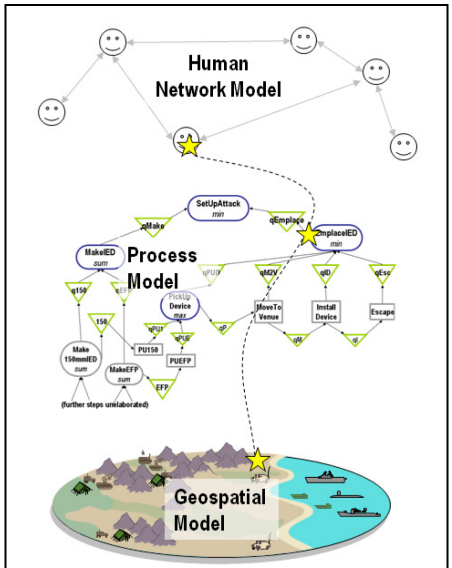

FIGURE 1.2 Inter-Tier Entities (ITES) coordinate the dynamics of multiple environments

Polyagents can explore multiple futures in a process graph just as they can in a geospatial map. In the example of FIGURE 1.3, some ghosts will sample one "Move" task, others will sample the other. This search can be influenced by the success of their venture by having the ghosts retrace their steps through process space, marking each task node with digital pheromone indicating the value received, and weighting the choices of subsequent ghosts by this value pheromone. The result will be the emergence of the most profitable route.

Another paper (Parunak, Brueckner et al. 2010) walks through our multi-perspective approach in more detail, and (Parunak, Sauter et al. 2009) describes its application to a military problem.

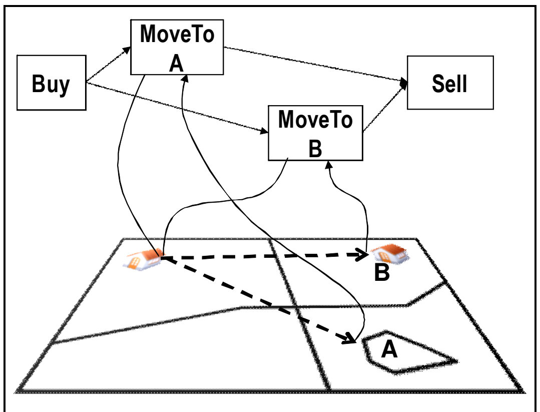

FIGURE 1.3 Inter-Tier Entities concurrently explore multiple perspectives

SOLVING THE PROBLEM

This architecture addresses the three challenges of multi-perspective modeling.

Combinatorics.—Ghosts interact by means of pheromones in whatever space they explore. Thus each ghost explores the interactions of its particular future with the distribution of futures of other entities, providing the same efficient exploration of multiple futures that we achieved in geospatial modeling..

Nonlinearities.—By actually emulating agent activities in the appropriate space, we obtain a much more realistic distribution of outcomes than if we sampled artificial distributions. In our example, the agents executing moves in physical space are subject to all the nonlinearities of interaction with other agents, and their completion times reflect a realistic distribution of possible experiences.

Model Maintenance.—The various subspaces can be constructed and tested separately, then combined at run-time as needed.

ANALYZING POLYAGENT ARTIFACTS

THE PROBLEM

Simulation models are excellent stand-alone tools for analyzing complex situations. However, they can be difficult to interface with other technologies. One particularly powerful set of methods is based on Bayesian probability theory. These methods require distributions over the alternatives that they explore. Conventional applications sample from distributions that are computationally convenient (e.g., the use of conjugate priors), but not necessarily realistic. A simulation model generates realistic distributions. We would like to derive statistical artifacts from our simulations to feed other, statistically-based, reasoning tools.

PHEROMONES ARE PROBABILITY FIELDS

Up to a normalizing constant, the pheromone field incremented by the ghosts of a given entity on a given page of the book of maps is a probability field estimating the entity's location at the time represented by that page over all possible futures explored by that entity's ghosts. This claim is supported by the dynamics of pheromone field strength. The strength of the field in a cell is augmented by a constant deposit D each time a ghost visits the cell, and decremented by a constant fraction E each time step. The strength $\varphi(t)$ for a single cell with a single ghost has dynamics.

$$\varphi(\tau) = E\varphi(\tau-1) + D = D\sum_{t=0}^{\tau-1} E^t = D\frac{1-E^t}{1-E} \tag{3}$$

In the continuous time limit, $\varphi(t)$ converges exponentially to $\frac{D}{1-E}$. This result has two important consequences.

First, the field converges if enough shifts of ghosts visit a page. We determine experimentally how many shifts are needed in each application.

Second, the evaporation rate E does not change over time. Thus the converged strength of a pheromone field is proportional to the amount of deposit, even in the presence of evaporation. If multiple ghosts visit a cell over time and deposit the same pheromone flavor, the converged strength of the field in the cell is proportional to the average number of deposits experienced by the cell per unit time. In other words, pheromone strength measures ghost traffic through a cell.

To compute the appropriate normalization, observe that all ghosts representing an avatar must pass through some cell on a given page as they move through the time interval represented by the page. The proportion of the ghosts that visit a given cell is equal to the ratio between their pheromone in that cell, and the total amount of pheromone deposited on the entire page. But this ratio is just the probability that the avatar will visit that cell.

We can use the field to estimate the probability that the entity is in a given region of the page. Let A be the total amount of the entity's pheromone on the entire page, and B the amount in a region of interest. Then B/A estimates the probability that the entity is in the limited region. The entity's most likely location is given by the center of mass of the probability field.

Often we are interested not in the location of an individual entity, but in the distribution of a group of entities (for example, all members of a team). In this case, all ghosts of entities on the same team increment the same field, which now reflects the probability of encountering any team entity over the area of interest.

In interpreting these fields, we must understand that the ghosts are moving under the constraints of a range of environmental influences on earlier pages (represented as pheromones of other flavors). The probability field that they estimate is thus not P(Avatar at location (x, t)), but

$$P(\text{Avatar at location } (x, t) | \text{Other conditions at } t' < t)$$

That is, the movements of individual ghosts are not independent of one another. They are all subject to the same conditions. However, those conditions form a Markov blanket for the locations of the ghosts, so *given* those conditions, the ghosts' locations (and thus the pheromone field that they generate) *can* be treated as independent samples of the avatar's location in space-time.

Elsewhere (Parunak 2009), we describe in more detail how to generate useful distributions, not only from the pheromone fields that ghosts lay down as they execute, but also from efficient logs of ghost state that we collect as they move from page to page.

CONCLUSION

Modern analysis requires models that can deliver statistically meaningful results from complex scenarios in a form that can interact with other analytic tools. The technologies that we have outlined in this paper address this need.

- By considering multiple futures concurrently, they avoid the sampling error of considering only one or a few trajectories.

- By modeling the location of entities in multiple environments concurrently and managing the flow of execution, they can configure complex environments easily from a library of standard components.
- The computational artifacts they produce can be interpreted readily by other systems.

REFERENCES

Erdös, P. and Rényi, A. (1960), "On the Evolution of Random Graphs." *Magyar Tud. Akad. Mat. Kutató Int. Közl.* 5: 17-61.

Gilmer, J.B., Jr. and Sullivan, F.J. (1998), "Alternative implementations of multitrajectory simulation." *the 30th Winter Simulation Conference*, Washington, D.C., United States, IEEE Computer Society Press.

Gilmer, J.B., Jr. and Sullivan, F.J. (2000), "Recursive simulation to aid models of decision making." *the 32nd Winter Simulation Conference*, Orlando, Florida, Society for Computer Simulation International.

Gilmer, J.B., Jr. and Sullivan, F.J. (2001), "Study of an ergodicity pitfall in multitrajectory simulation." *the 33nd Winter Simulation Conference*, Arlington, Virginia, IEEE Computer Society.

Parunak, H.V.D. (2007), "Real-Time Agent Characterization and Prediction." *International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'07), Industrial Track*, Honolulu, Hawaii, ACM.

Parunak, H.V.D. (2009), "Interpreting Digital Pheromones as Probability Fields." *the 2009 Winter Simulation Conference*, Austin, TX.

Parunak, H.V.D., Belding, T.C. and Brueckner, S. (2007), "Prediction Horizons in Polyagent Models." *Sixth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS07)*, Honolulu, HI.

Parunak, H.V.D. and Brueckner, S. (2006), "Concurrent Modeling of Alternative Worlds with Polyagents." *the Seventh International Workshop on Multi-Agent-Based Simulation (MABS06, at AAMAS06)*, Hakodate, Japan, Springer.

Parunak, H.V.D., Brueckner, S.A. and Bisson, R. (2010), "Agent Interaction, Multiple Perspectives, and Swarming Simulation." *the International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2010)*, Toronto, Canada, IFAAMAS.

Parunak, H.V.D., Sauter, J. and Crossman, J. (2009), "Multi-Layer Simulation for Analyzing IED Threats." *IEEE International Conference on Technologies for Homeland Security (HST 2009)*, Waltham, MA, IEEE.