Generation and Analysis of Multiple Futures with Swarming Agents (Extended Abstract)

H. Van Dyke Parunak
Vector Research Center of TTGSI
3520 Green Court, Suite 250
Ann Arbor, MI 48105
+1 734 302 4684
van.parunak@newvectors.net

ABSTRACT

Most agent-based modeling techniques generate one trajectory per run, greatly under-sampling the space of possible trajectories. Swarming agents that interact through digital pheromones can explore many alternative futures in parallel, based on an interpretation of pheromone fields as probability fields that yields more information from them than swarming models usually yield, and also facilitates integration with probability-based AI mechanisms.

Categories and Subject Descriptors

I.2.11 [**Computing Methodologies**]: Distributed Artificial Intelligence – *multiagent systems.*

General Terms

Algorithms, Measurement, Experimentation, Theory.

Keywords

Polyagent, swarming, probability, agent-based modeling.

1 INTRODUCTION

Agent-based models offer important benefits over other modeling techniques such as equation-based models (EBMs) [5, 6]. However, their relatively slow execution is a disadvantage in highly nonlinear domains, since it is costly to sample the space of alternative futures that such a domain supports.

Polyagents [4] use multiple swarming agents to represent each domain entity. Each agent explores a different future of its entity concurrently. When these agents interact, not directly but through fields that they deposit in the environment, they explore a space of alternative futures far greater than the number of agents per entity. If the system is modeling e entities with g agents per entity, the number of futures sampled is on the order of g^e. This sampling of alternative futures can yield probability distributions over a variety of propositions of interest, and permits integration with probability-based machine reasoning systems.

We discuss the problem of multiple futures (Section 2) summarize the polyagent modeling construct (Section 3), show how pheromones can be interpreted as probabilities (Section 4), and estimate the number of futures such a system explores (Section 5).

2 MULTIPLE FUTURES

Successive runs of an agent-based model often yield different trajectories, because of imperfect knowledge of the agents' internal states or details of the environment, non-determinism in agent decisions, or iterated nonlinearity in decisions or interactions that lead to chaotic dynamics. During a run of length τ, each of n entities will experience one of as many as n^{τ} possible histories [4]. Even multiple runs with different random seeds sample a vanishingly small portion of the space of possible entity histories.

In equation-based modeling, stochastic differential equations can propagate uncertainty through the model. However, such approaches are subject to the recognized weaknesses of EBM's resulting from their use of aggregated population characteristics.

One approach to multi-trajectory simulation [1] evaluates possible outcomes at each branch point stochastically and propagates the most likely. This selection is required by the high cost of following multiple paths, but avoiding low-probability paths violates the model's ergodicity and compromises accuracy [2].

3 THE POLYAGENT CONSTRUCT

The polyagent modeling construct [4] represents each domain entity as a persistent *avatar* that manages a stream of transient *ghosts*, each of which explores an alternative future for the entity in a simulated world. As the ghosts of different avatars interact, each one explores an alternative future for its entity, complete with the full range of possible interactions that might result from the alternative futures of other entities. 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 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, specific to each avatar, are deposited by the ghosts as they move about, and evaporate over time to discard obsolete information. An avatar releases its ghosts in successive shifts so that they can respond to pheromones deposited by other ghosts. A ghost's behavior is determined by a weighted sum of the pheromones it senses in its vicinity.

4 PHEROMONES ARE PROBABILITIES

Up to a normalizing constant, the pheromone field incremented by the ghosts of a given entity on a given page of a 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.

Cite as: Generation and Analysis of Multiple Futures with Swarming Agents (Extended Abstract), H.V.D. Parunak, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 1549-1550

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

In the continuous 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. 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 the map. 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.

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(Avatar at location (x, t)|Other conditions at $t' < t$)

That is, movements of individual ghosts are not independent. 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 the pheromone field that they generate) *can* be treated as independent samples of the avatar's location in space-time.

Ref [3] shows how these fields can be used to derive a variety of useful distributions over a simulation.

5 COUNTING POSSIBLE FUTURES

In the real world, one entity's behavior can depend on the presence or absence of another entity. A ghost senses other entities through their fields, and thus samples all of the locations of those entities that their ghosts have explored. How many different futures does this approach explore?

Over σ shifts, 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 representing 60 successive future time steps.

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. A naïve estimate of the number of possible futures as $(g*\sigma)^n \sim 3.2*10^{11}$. This is an overestimate:

1. The number of ghosts that have visited a given page depends on the page. Pages further in the future see fewer ghosts.
2. A ghost interacts with later ghosts through its field, which evaporates over time. So not all ghosts count 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 $g\sum_{i=0}^{\sigma-\tau} E^i = g\frac{1-E^{\sigma-\tau}}{1-E}$. 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 x 10^7, still far more than a single-trajectory simulation can explore.

6 CONCLUSION

Digital pheromones are an effective means of coordination among multiple agents in a wide array of domains. In many systems, they summarize the behaviors of many similar agents, so it makes sense to interpret them as probability fields over that space of behavior. This interpretation has two benefits. First, it justifies viewing the swarming search as concurrent exploration of many possible futures. Second, it can estimate dynamic, nonstationary distributions for other algorithms, distributions that would otherwise be extremely difficult to compute.

Polyagent-based exploration of multiple futures achieves its breadth at the expense of detail. When ghosts interact through probability fields, some information that would be gained in a one-on-one interaction is lost. Nevertheless, this approach offers access to distributions that would otherwise be inaccessible.

This work benefited from experiments by Bob Matthews, Jorge Goic, and Andrew Yinger, and comments by John Sauter. Jacob Crossman, and a referee. A paper with more detail is available at www.newvectors.net/staff/parunakv/AAMAS2010MultiFuture.pdf.

7 REFERENCES

[1] J. B. Gilmer, Jr. and F. J. Sullivan. Alternative implementations of multitrajectory simulation. In *Proceedings of the 30th Winter Simulation Conference*, pages 865-872, IEEE Computer Society Press, 1998.

[2] J. B. Gilmer, Jr. and F. J. Sullivan. Study of an ergodicity pitfall in multitrajectory simulation. In *Proceedings of the 33nd Winter Simulation Conference*, pages 797-804, IEEE Computer Society, 2001.

[3] H. V. D. Parunak. Interpreting Digital Pheromones as Probability Fields. In *Proceedings of the 2009 Winter Simulation Conference*, pages 1059-1068, 2009.

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

[5] N. M. Shnerb, Y. Louzoun, E. Bettelheim, and S. Solomon. The importance of being discrete: Life always wins on the surface. *Proc. Natl. Acad. Sci. USA*, 97(19 (September 12)):10322-10324, 2000.

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