Presented at the Workshop on Evolutionary Computation and Multi-Agent Systems (ECOMAS 2001), GECCO 2001, San Francisco, CA, pp 321-324.

Tuning Synthetic Pheromones With Evolutionary Computing

John Sauter
ERIM
PO Box 134001
Ann Arbor, MI 48113-4001
+1 734 623 2513
jsauter@erim.org

H. Van Dyke Parunak
ERIM
+1 734 623 2509
vparunak@erim.org

Sven Brueckner
ERIM
+1 734 623 2529
sbrueckner@erim.org

Robert Matthews
ERIM
+1 734 623 2510
rmatthews@erim.org

Abstract

Agents guided by synthetic pheromones can imitate the dynamics of insects. These systems are well suited to problems such as the control of unmanned robotic vehicles. We have developed a model for controlling robotic vehicles in combat missions using synthetic pheromones. In the course of our experimentation, we have identified the need for proper tuning of the algorithms to get acceptable performance. We describe pheromones in natural and synthetic systems, and describe the mechanisms we have developed. The role of evolutionary computing in offline and online tuning is discussed.

1 PHEROMONE FIELDS

From an engineering perspective, pheromones are a particularly attractive way to construct a potential field that can guide agent-based systems. Pheromones have been used in material transport routing (Parunak, 1987), combinatorial optimization (Bonabeau, 1999), and factory control (Brueckner, 2000). This paper discusses the use of pheromone fields in constructing paths for movement of robotic vehicles.

1.1 NATURAL SYSTEMS

Many social insect species perform impressive feats of coordination without direct inter-agent coordination or complex reasoning. Instead, they deposit and sense pheromones (chemical scent markers) in the environment (Parunak, 1997). For example, ants construct networks of paths that connect their nests with available food sources. Mathematically, these networks form minimum spanning trees (Goss, 1989), 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 nest pheromones while dropping food pheromones if they are holding food.

Figure 1a shows a simulated pheromone field deposited by a swarm of ants wandering out from their nest in search of food. Initially, the field is roughly circularly symmetrical, and serves to guide food-bearing ants back home. Once some ants find the food and begin returning home, the field rapidly collapses into a path (Figure 1b).

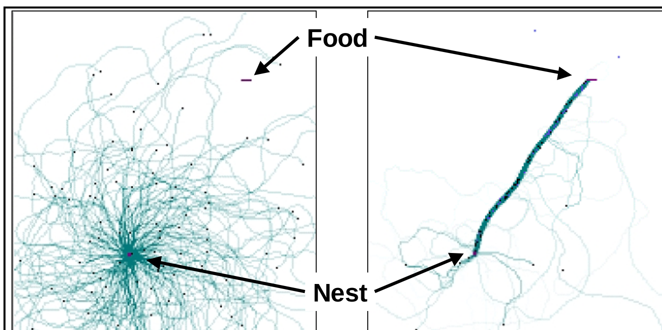

a. Initial *Field* **b. Final *Path***

Figure 1: Path Condensation

The real world provides three operations on chemical pheromones that support purposive insect actions. It *aggregates* deposits from individual agents, providing fusion of information across agents and through time. It *evaporates* pheromones over time removing old or inconsistent information. It *diffuses* pheromones to nearby places, disseminating information locally. Under continuous reinforcement, the pheromone field strength stabilizes rapidly, as a concave function of time. Thus new information is quickly integrated into the field, while obsolete information is automatically forgotten through pheromone evaporation.

The pheromone field constructed by the ants in the environment is in fact a potential field that guides their movements. Unlike many applications of potential fields, both ant behavior and pheromone field maintenance are decentralized. Ants interact only with the pheromones in their immediate vicinity, by making deposits and reading the local strength of the pheromone field. The nature of pheromone propagation and evaporation provides a naturally distributed mechanism for building and maintaining these fields.

1.2 SYNTHETIC PHEROMONES

In DARPA's JFACC program (Parunak, 2000) we used a synthetic pheromone field to guide unmanned robotic vehicles (URV's) through the battlespace (Figure 2). In this scenario, robotic vehicles seek to find and destroy targets, which are defended. The URVs are guided by potential gradients centered on the targets while avoiding gradients centered on the threats.

An implementation of synthetic pheromones has two components: a network of *place agents* that maintain the pheromone field and perform aggregation, evaporation, and diffusion, and the *walkers* which deposit and react to the field maintained by the environment. In JFACC, we tile the physical space with hexagons, each representing a place agent with six neighbors. The underlying mathematics of the pheromone field, including critical stability theorems, is described in (Brueckner, 2000).

A walker agent is associated with one place agent at any given time. It can read the current strength of pheromones at that place and at each of its neighbors, and can deposit pheromones into the place. A walker moves from one place to another by spinning a roulette wheel whose segments are weighted according to this set of strengths.

1.3 GHOST AGENTS

The path emergence illustrated in Figure 1 is the result of the interactions among many walkers performing a stochastic, Monte Carlo search of their environment. In real applications, it is not be feasible to have URVs explore the domain in this manner.

In our implementation a URV sends out many unembodied walkers, which we call *Ghosts*. Ghosts can move as fast as the network among place agents can carry them. As they sense pheromones in the environment, they deposit their own flavors of pheromones, which condense into paths like that in Figure 1. The URV then follows this path. As the URV moves, it continuously sends out ghosts, so that the path it follows is being constantly revised to accommodate dynamic changes in the environment.

2 ENGINEERING GHOSTS

We have explored several basic mechanisms to guide the path formation capabilities of ghosts. We will present one of these mechanisms here: using combinations of multiple pheromones. Due to the large number of tunable parameters this mechanism is an ideal candidate for evolutionary computing.

2.1 PHEROMONE VOCABULARY

Different pheromone flavors may reflect different features of the environment (e.g. hostile (Red) and friendly (Blue) units). In our applications, the ghost's choice function weights the various input pheromones to create a single "net pheromone" for each neighbor that is used in weighting the roulette wheel for determining the next move. The basic pheromone flavors are:

- *RTarget*: emitted by a target (e.g. Red headquarters).
- *GTarget*: emitted by a ghost who has encountered a target and is returning to the URV.
- *GNest*: emitted by a ghost who has left the URV and is seeking a target.
- *RThreat*: emitted by a threat (e.g., Red air defense)

In addition, the ghost may know *Dist*, an estimate of the distance of the target when the target location is known.

We experimented with several different forms of the equation. Manual manipulation of the equation yielded the current form:

$$\frac{\theta \cdot RTarget + \gamma \cdot GTarget + \beta}{(\rho \cdot GNest + \beta)(Dist + \varphi)^{(\delta + \alpha(RThreat+1))} + \beta},$$

3 TUNING THE EQUATION

Our experiments show this path formation dynamic to be extremely robust and adaptive. Figure 2 shows the formation of a path from a friendly airbase (lower right) to the nearer (in terms of safest path) of two targets (the house-shaped icons) with equal strength, avoiding threats (the radar icons). If we increase the strength of the left-hand target to twice that of the closer target, the path will lead there instead.

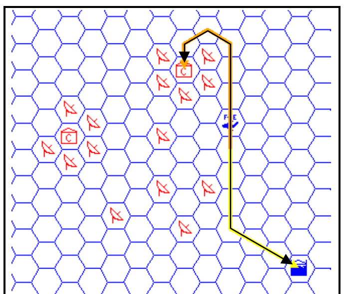

Figure 2: Path to the nearer of two targets.

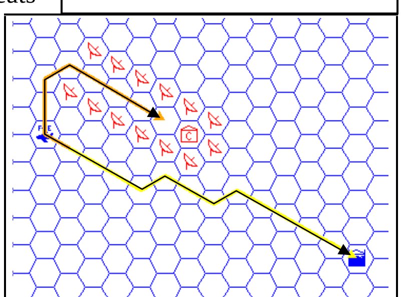

Figure 3: Threading a gauntlet.

Figure 3 shows the formation of a path to a target protected by a long narrow gauntlet of threats oriented away from the base, a configuration that resists classical potential field methods.

We also experimented with sealing the open end of the gauntlet (so the ghosts must eventually decide to pierce the air defense shield to reach the target), and with placing the target at a great distance from the base. Table 1 shows the manually selected parameters that resulted in good path formation.

Table 1: Manually Selected Parameter Settings

Parameter	Two Targets	Gauntlet	Sealed Gauntlet	Distant Target
α	10	10	**3**	10
δ	0	0	**4**	**9**
φ	1.5	1.5	1.5	**4**
ρ	1	1	**3**	1
θ	1	**0**	1	**0**

3.1 OPPORTUNITIES FOR EVOLUTIONARY COMPUTING

Our experience in working with these algorithms strongly suggests that techniques from evolutionary computing would be ideal for tuning. The process of experimenting with several different forms of the equation and different parameter settings was long and involved. Evolutionary Computing (EC) provides an attractive means to formulate and tune the equations. Despite the seeming benefit from combining EC with synthetic pheromones, most of the literature has focused on contrasting ant optimization algorithms and EC. Perhaps this is due to the rather simple forms most ant optimization algorithms have taken to date (Coren, 1999). Some issues that EC could address:

1. *Evolving the form of the equation* - A parse tree could represent the equation with the operators from the set [+,−,×,÷,^].

2. *Offline tuning for a wide range of scenarios* - this would avoid the need to develop new parameter settings for each new scenario.

3. *Higher level control input* - There are certain non-linearities in the pheromone behavior that make it difficult to adjust the parameters manually for desired outcomes. For example, as the relative distance to two diametrically opposed targets increases, the strength of the farther target must be increased exponentially to attract ghosts. This makes it difficult for commanders to set target priorities by simply adjusting the strengths. EC could evolve the correct parameter settings to get the desired behavior.

4. *Dynamic adjustment* - Even with the best offline tuning there will be situations where new behaviors must be developed during the operation. EC techniques could provide a mechanism to dynamically adjust the behavior during operation.

3.2 OFFLINE TUNING WITH GA

Genetic Algorithms (GAs) could be used to find the best set of parameters for path formation under a number of different scenarios. The typical range on the tuning parameters is small. A straightforward bit encoding of the parameters should suffice for the GA. The initial population could be chosen randomly.

Developing an appropriate fitness function for the static path formation problem is straightforward. A single path from a stationary source to a single stationary target can be evaluated based on the following criteria:

Length of the path (L_p). Shorter paths are better unless they represent unacceptable risk.

Risk of the path (R_p). This represents the relative risk from enemy threats encountered along the path.

Time to form the path (t_p). The time it takes for the ghosts to form a path to the target.

Multiple configurations of target and threat locations would be used in the evaluation. The tradeoff among these metrics is not known so evaluation would use a Pareto multiobjective optimization approach.

3.3 TUNING DYNAMIC PATH FORMATION

There is the need to evaluate the ability of the ghosts to perform continuous path planning as the environment changes. In a real environment the URV, targets and threats are all moving. Additionally threats that were previously hidden may “pop-up”. Evaluating the performance of the ghosts in such a dynamic environment is more difficult. Even if suitable metrics could be developed, a single evaluation of one individual could take several minutes of computer time. Identifying suitable parameters could take over a day of compute time. Since the environmental dynamics are varying on the order of minutes, such a delay in reaching a decision is unacceptable.

In addition to the evaluation delays, GAs tend to suffer from premature convergence in the face of dynamic problems. Some of the techniques that have been used to overcome this problem are:

1. Identify changes in the environment and then increase diversity by means of increased mutation.

2. Maintain diversity by including new random individuals in the population in every generation.

3. Supply the EA with a memory, e.g. by using diploidy or an explicit memory, so that the EA can recall useful information from past generations (Goldberg, 1987).

4. Use co-evolution with tournaments for evaluation or multiple populations from which individuals are selected for evaluation of group fitness.

We believe that co-evolution provides the most promising approach to maintaining the best ghost behaviors in a dynamic environment. Ghosts that wander too close to threats would die in simulated combat. Ghosts that find high strength targets would be rewarded with additional lifetime in order to have time to find their way back home. The ghosts returning to the URV from the best targets are selected for the next round of breeding. The ghosts that die in the environment never return and therefore do not participate in the breeding.

Using co-evolution, symbiotic relationships could develop that improve the overall fitness of the population. For example, risk-taking ghosts (low α) would be willing to penetrate air defenses in the hopes of finding targets. They cooperate with more risk-averse ghosts (high α) to create the shortest-safest path to the target.

Some preliminary experiments were performed using simply evolutionary strategies. Ghosts were created with random values (uniform distribution about a mean) for each parameter. A Ghost's life was reduced by a factor of two when it chose to move into a location with a threat. Ghosts that returned after finding a target were placed in a queue for the remainder of their life. While there were two or more ghosts in the queue, new ghosts were created by randomly choosing a set of parameters from two of the ghosts in the queue. Parents were either chosen from the head of the queue only (labeled "queue") or from a round robin of all the ghosts in the queue ("round robin"). Figure 4 shows the performance of these two strategies versus a random population and a fixed population of the best hand-tuned set of parameters. The plot shows the strength of the *GTarget* pheromone in the base neighbor cell that is closest to the target. The evolutionary strategies were able to generate a population of ghosts that did significantly better at building a strong path to the target in a shorter

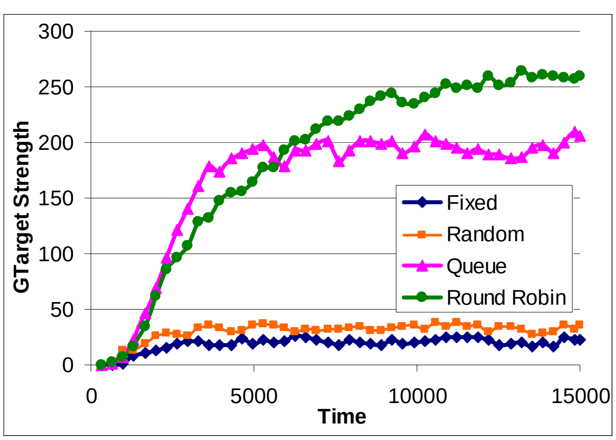

Figure 4: Performance of Simple Evolutionary Strategies

period of time. By starting with median values for all parameters, this technique was able to solve all the problems listed in Table 1 without any hand tuning. These simple experiments show the potential of these algorithms for tuning pheromone systems.

4 CONCLUSION

Synthetic pheromones are a powerful mechanism for controlling the movement of agents through space. They provide the elegance of potential field methods, with support for integrating diverse information sources, processing information in a completely distributed environment, and coping with dynamic changes in the landscape. EC can play an important role in tuning the behavior of these systems in a dynamic environment.

Acknowledgements

This work was supported in part by the DARPA JFACC program. Our research team included O. Gilmore, M. Nandula., J. Posdamer, E. Feibush, and E. Greene.

References

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

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

D. Coren, M. Dorigo, and F. Glover, New Ideas in Optimization. McGraw Hill, 1999.

D. E. Goldberg and R. E. Smith. Nonstationary Function Optimization Using Genetic Algorithms with Dominance and Diploidy. In J.J. Grefenstette, editor, *2nd Intl. Conf. on Genetic Algorithms*, pages 59-68. Lawrence Erlbaum Associates, 1987.

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

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

H. V. D. Parunak. Adaptive control of Distributed Agents through Pheromone Techniques and Interactive Visualization. 2000., www.erim.org/cec/projects/adaptiv/.

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, AAAI, 1987.28, 1987.