

Evolving Cooperative Pheromone Usage in Digital Organisms

Brian D. Connelly, Philip K. McKinley, and Benjamin E. Beckmann
Department of Computer Science and Engineering
Michigan State University
East Lansing, Michigan 48824
{*connel42,mckinley,beckma24*}@cse.msu.edu

Abstract—The use of chemicals to communicate among organisms has enabled countless species, from microorganisms, to colonies of insects, to mammals, to survive and flourish in their respective environments. Ants, arguably nature’s most successful exploiters of this behavior, have evolved the use of pheromones to communicate in a wide range of situations, including mating, colony recognition, territory marking, and recruitment to new nest sites and food sources. We examine the evolution of the use of pheromones to aid in the location of, and migration to, a target area by groups of digital organisms. In an initial set of experiments, these organisms evolved efficient patterns of exploration that obviated the need for pheromones. When evolved in a more adverse environment, organisms again evolved effective search strategies, but also evolved the use of pheromones to enable the task to be completed by group members more quickly and with fewer movements. We also show that evolved organisms are more robust and better able to react to a change in the environment than a hand-built solution. This work demonstrates the complexities that exist in the evolution of pheromone-enabled cooperation and provides insight into the behaviors executed by seemingly simple organisms in nature.

I. INTRODUCTION

How multiple agents communicate to coordinate and make effective group-level decisions is a fundamental question that spans many areas of research. In terms of evolution, building even simple group behaviors is very difficult, as there is a cost associated with cooperation, which individual-level selection would remove over time unless it produced a benefit that outweighed the cost. Indeed, when the success of the group is considered instead of that of an individual, cooperative behaviors can be a major contributing factor. Stigmergy, the indirect communication through traces left in the environment, is a particular method of communication that has evolved and is used with great success in nature by a variety of organisms, including many species of microorganism [1], insects [2], [3], marine organisms [4], and mammals [5]. Highly-developed eusocial insects such as ants produce pheromones [6], which can trigger certain behaviors in other members of their species. Ants have been observed to use pheromone trails to navigate toward food sources and potential nest sites hundreds of meters away [7].

Different pheromones and pheromone blends can be used to attract mates, signal warnings, and mark territories, among several other behaviors. By leaving traces in their environment, organisms are able to effectively communicate without the need for other complex or long-distance mechanisms. Also, the completely decentralized nature of pheromone usage makes this method robust in the face of changing environments.

The goal of the study described here is to provide a better understanding of pheromone communication by focusing on how populations can *evolve* the use of pheromone trails to help members of the group find a target, a task similar to the mass recruitment of ants to a food source [7]. By allowing organisms to mark their environment, other organisms in that population are able to be recruited in less time and with fewer movements. To carry out this study, we extended the Avida digital evolution platform [8] to support several pheromone-related primitives. During initial runs, organisms evolved an effective search strategy that resulted in their locating the target without the need for pheromones. Based on these evolved organisms, we hand crafted an organism that followed a similar search strategy, but also used pheromones to help others once it had found the target. Although these manually-constructed organisms were successful and retained their use of pheromones, the solution proved to be brittle, and did not translate well into different environments. When exposed to a more adverse environment, organisms again evolved effective search strategies, but also evolved the use of pheromones, enabling efficient task completion. The main contribution of this work is to provide insight into the factors that led to these cooperative behaviors.

The remainder of this paper is organized as follows. Section II provides a brief overview of pheromone-related research in several different fields. Section III describes the Avida digital evolution platform, and Section IV highlights the additions made to Avida to support this study. Experimental results and analyses are presented in Section V, while conclusions are drawn and future avenues for extending this work are posed in Section VI.

II. RELATED WORK

How groups of organisms develop cooperation and communication is a very complex problem that has been studied extensively across many disciplines. Typically, an individual pays a cost to cooperate with others, who may reap some benefit as a result. Because the contributing organism may not directly receive the benefits of its cooperation, selection based strictly on the success of the individual would suppress such behaviors. Selection at the group level, however, can be strong enough to preserve cooperative traits that benefit a group of organisms, and evolved cooperation has been observed in numerous organisms as simple as bacteria [1]. The evolution of cooperation has been studied in both natural and artificial organisms. This section highlights some of the work conducted in these areas.

Evolving Cooperation. Several conditions have been observed that indicate whether or not group selection will favor cooperation [9]. Nowak summarized the mechanisms involved in the evolution of cooperation—kin selection, direct reciprocity, indirect reciprocity, network reciprocity, and group selection—and defined the conditions under which natural selection favors them [10]. Wilson has argued that individual and kin selection can even be dissolutive to group behavior in certain cases, and that group selection is the dominant force driving evolution towards more advanced cooperation [11].

Pheromone Evolution. An area of particular interest is the evolution of pheromone diversity [12]. This problem is important in sex pheromones, where cross-attraction between species can have strong negative effects. The evolution of the biomechanics required to produce and sense pheromone molecules has also been examined [13]–[17], as has the use of pheromones to regulate recruitment [18]. Typically, the majority of colony members choose the most profitable among multiple food sources. This can be seen as a natural byproduct of the use of pheromones—lesser-used trails will not be as heavily re-enforced and will consequently be chosen less frequently. Based on observed behaviors, several models have been developed to capture recruitment behavior in the presence of multiple targets [19].

Pheromone Use in Evolutionary Computation. The evolution of pheromone use has been studied in evolutionary computation (EC) environments since the early 1990s. These works tend to focus on harnessing pheromones to be used as navigational cues in robotics and other multi-agent systems [20]. Ant-colony optimization [21] has been applied successfully to several computational tasks such as network routing [22] and the traveling salesman problem [22]. In the AntFarm project [23], an early study of EC-based pheromones, a genetic algorithm was designed to closely model the real-life capabilities of ants. Homogeneous groups of organisms were rewarded

based on the amount of food brought from a source to the nest. In addition to the ability to sense and produce pheromones, organisms were given a virtual compass to aid in navigation. Experiments showed an increase in the amount of food brought back to the nest over time.

More recently, Panait and Luke [24] used genetic programming to evolve both foraging behavior and effective pheromone levels for foraging. Their task involved organisms finding a food source and returning to the nest carrying some of that food. The approach used two different pheromones—one for marking the path to the food source and another for marking the path to the nest. Evolved organisms used an exploration strategy where they moved toward regions with more food pheromones and also where less nest pheromone was present. Further studies found that the ability to mark cells with the maximum amount of pheromone possible at that location resulted in significantly more food delivered to the nest than with fixed increments in pheromone.

In related work, Sauter *et al.* [25] used genetic algorithms to tune agents' use of, and attractions to, four different pheromones used to control robotic vehicles in military scenarios. Floreano *et al.* [26] evolved software controllers for robots to use light to signal the presence of a food source or a poison source. That work explored strategies evolved in different selection regimes and levels of population diversity. It was found that cooperative communication occurred primarily when a population consisted of closely-related organisms (in their case, identical) and when selection occurs at the group level. This finding supports our use of *demes* in Avida, where homogeneous subpopulations compete against one another in performing group tasks.

III. THE AVIDA DIGITAL EVOLUTION PLATFORM

Avida is a software platform used to study the evolution of self-replicating computer programs in a fixed-size environment [8]. These computer programs, or “digital organisms,” compete for CPU cycles by completing *tasks*, further discussed below, in an environment configured by the user. Avida has been successfully used to explore fundamental evolutionary processes [27], the emergence of adaptive behaviors [28], and the evolution of group-based cooperation and communication [29]. Avida and other EC tools have been increasingly used to solve complex problems in engineering [30].

As depicted in Figure 1, each Avida organism, or *Avidian*, exists independently of other organisms in its environment and is placed within its own *cell* in that environment. An organism's behavior is defined by a circular set of assembly-like instructions (its “genome”), which execute sequentially on virtual hardware allocated to that organism. In the study described here, the virtual hardware comprises a CPU, three 32-bit registers, and

two stacks that store up to ten 32-bit numbers. The CPU has four *heads*, which serve as pointers to locations in the genome. These heads are used to determine which of an organism’s instructions to execute, enable jumps and loops in the execution of an organism’s genome, and aid in self-replication. Input- and output buffers also exist in the virtual hardware, and values written to and read from these buffers are commonly used to determine if certain tasks have been completed.

In the beginning of a run, the population is seeded with an ancestral organism. Typically, this organism’s genome contains only the instructions necessary to replicate. All other behaviors exhibited by that organism’s descendants must be evolved. Avidians are self-replicating, and offspring are typically placed in a random neighboring cell.

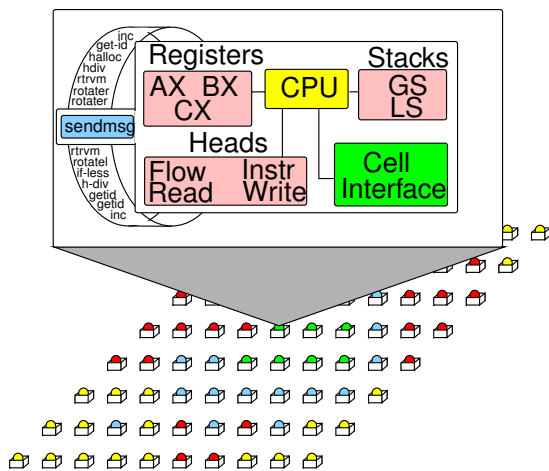


Fig. 1. The structure of an Avidian (top), and a population containing multiple genomes (bottom) [31].

Avidians are able to move about their environment [32]. Each organism is born facing a particular neighbor cell. The *rotate-left* and *rotate-right* instructions allow an organism to change its facing to the next cell to the left or right, respectively. The *rotate-label* instruction allows an organism to change its facing to a direction indicated by subsequent no-operation instructions. Organisms move into the cell currently faced through the use of the *move* instruction. If another organism is currently in the faced cell, the positions of two organisms will be swapped.

In Avida’s energy model [28], an organism’s metabolic rate is determined by its current energy level. An organism with a higher metabolic rate will probabilistically execute more instructions than an organism with a lower metabolic rate. A side effect is that proportionately more energy is used per instruction. The unit of time in Avida is the *update*. By default, organisms are configured to execute 30 instructions per update on average; however, this amount depends on the organism’s metabolic rate.

Organisms can gain energy through the successful completion of *tasks*, which are defined by the user in terms of an organism’s observable behaviors (its *phenotype*). For example, a task might be to perform a mathematical or logical operation [27], mitigate an attack [33], or communicate with a neighboring cell to solve a distributed problem [29]. The Avida task mechanism creates selective pressures in the population, as successful organisms are able to replicate more quickly and spread throughout the population. Tasks may also involve the use of resources. By executing instructions, organisms are able to sense, consume, and produce resources. Resource levels in the environment can fluctuate over time through inflow, outflow, diffusion, or consumption.




Fig. 2. An Avida population divided into sixteen demes. Demes are isolated subpopulations, each capable of replication. When a deme replicates, it replaces a randomly-selected target deme [31].

Group-level selection links the survival of an individual to the survival of its group [34]. It therefore becomes beneficial to the individual to contribute to the success of its group. In Avida, the population of organisms can be divided into independent subpopulations, or *demes*, as shown in Figure 2. New organisms are born into the deme of their parent, and migration from deme to deme is not permitted. In this study, each deme has a germline [31], a shared source of genetic material. This results in a subpopulation of organisms with a common ancestral genome. Groups of highly-related individuals have been shown to be more successful in evolving cooperative strategies while also preventing the emergence of *cheaters*, organisms that take advantage of the benefits of cooperation but that do not contribute to it [26]. A deme’s organisms can cooperate to satisfy *predicates*, which allow the deme to replicate. As shown in Figure 2, when a deme replicates, another deme is randomly selected, and its organisms are replaced with organisms created from the deme’s germline. It is during this process that mutations occur in the replicating deme’s germline. Each mutation inserts, deletes, or modifies an instruction at a random point in the genome.

IV. METHODS

We extended Avida in two main ways to support this study. First, we defined several movement-

and pheromone-related instructions. Second, we created deme-level predicates that allow subpopulations of organisms to be replicated once certain behaviors have been exhibited. Here, we briefly overview these Avida extensions. Detailed information about the instructions, configuration variables, logging, and visualization tools can be found in an accompanying technical report [35].

A. New Instructions

Pheromone Instructions. A pheromone resource was created to support pheromone-related instructions. Initially, each cell contains zero pheromone. Organisms can sense the amount of pheromone at a cell as well as increase that value by “dropping” pheromone. Each organism has a boolean variable that represents its dropping status. When this status is **true**, the organism will drop pheromone with any movement instruction. The organism sets its pheromone-dropping status to **true** by executing the `phero-on` instruction. Dropping during subsequent movements can be turned off using the `phero-off` instruction. The amount dropped by an organism and its distribution are configured per run. The mode used for all runs in this work was to drop half of the pheromone into the source cell and half into the destination cell during each movement, effectively creating a pheromone trail between the two cells. An organism can sense the amount of pheromone in its own cell and the cell it is facing by invoking the `sense-pheromone` and `sense-pheromone-faced` instructions, respectively. Conditional instructions `if-pheromone` and `if-not-pheromone` execute the next instruction if the current cell does or does not contain pheromone.

Movement Instructions. Organisms are able to move throughout the environment using movement instructions. An organism moves into the cell faced by executing the `move` instruction; see Figure 3(a). As illustrated in Figure 3(b), the `exploit` instruction moves the organism into the neighboring cell with the highest pheromone level. When no pheromone is present, the organism simply moves forward. The `explore` instruction rotates the organism to a random facing and moves it into the faced neighbor cell, ignoring pheromone levels or neighboring target cells. The `movetarget` instruction, illustrated in Figure 3(c), moves the organism to a neighboring target cell. As with `exploit`, if no neighboring cell is part of a target, the organism is simply moved one cell forward.

Target Instructions. Analogous to `sense-pheromone` and `sense-pheromone-faced`, the `sense-target` and `sense-target-faced` instructions sense the presence of a target in the current or faced cell, respectively. The `if-target` and `if-not-target` instructions execute the next instruction if the current cell is a target cell or is not a target cell, respectively.

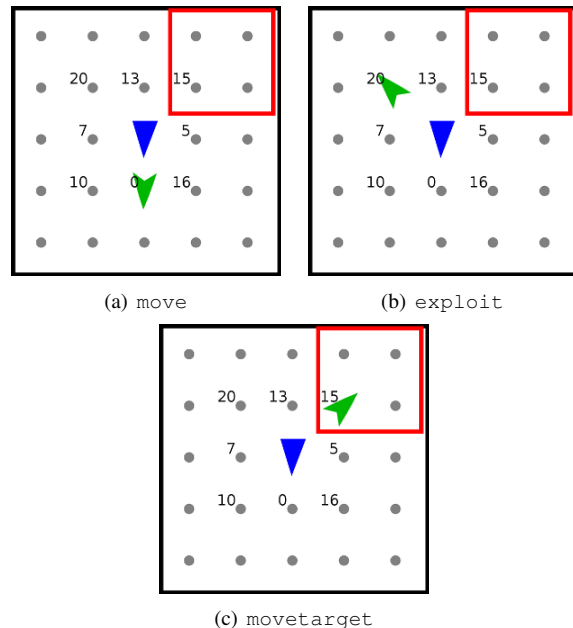


Fig. 3. Different Avida movement behaviors. The organism’s current position and facing are indicated with a triangle. Its new position and facing are indicated with a chevron. A 4-cell target exists in the upper right portion of the environment, and pheromone levels of neighboring cells are indicated in a cell’s upper left corner. (a) `move` simply moves the organism to the cell it is facing. (b) `exploit` moves the organism into the neighbor cell with the highest pheromone level. (c) `movetarget` moves the organism towards a neighbor marked as a target cell, ignoring the pheromone levels of neighboring cells.

B. Replication Predicates

Replication predicates are conditions that, when satisfied by a deme, allow that entire deme to replicate. Replicated demes continue to occupy the original deme and also start populating another randomly-selected deme. Acting as a group-level selection mechanism, faster-replicating demes are able to spread their genetic code to cover more of the global environment. It is during deme replication that mutations occur in the germline.

As the goal of this work is to investigate the evolution of the ability of a population to cooperatively find a target, a predicate was created that is satisfied when a prescribed number of unique organisms reaches any of the cells that make up an active target. A deme is replicated once it has satisfied this predicate. This behavior is reminiscent of quorum sensing in microbial communities: once a group has congregated or formed a quorum, the organisms change their behavior to work towards a new task [1].

V. EXPERIMENTS AND RESULTS

Several runs, each comprising 10 instances of Avida, were performed to investigate the conditions necessary for groups of organisms to evolve the cooperative use of pheromones in locating a target. Since each instance started with a different random seed, each followed its own evolutionary pathway. The main parameter values

TABLE I
PARAMETER VALUES USED DURING RUNS

Parameter	Value
Number of Demes	100
Organisms per Deme	10 or 20
Orgs required to target	8 or 16
Grid Size	40x40 or 20x20
Target Size	2x2 or 1x1
Germline Copy Mutation Rate	0.0075
Germline Insertion Mutation Rate	0.05
Germline Deletion Mutation Rate	0.05
Pheromone at Target	10 or 1000
Pheromone Dropped Per Move	10
Pheromone Evaporation Rate	0

used in these runs are shown in Table I. The environment was divided into 100 independent demes, each consisting of a toroidal grid of cells. At the beginning of a deme's existence, an organism is created from the germline and is inserted into a random cell. Organisms are injected periodically until the population reaches a configured size. Demes were allowed to replicate when a specified number of unique organisms had reached the target cell, or when the age of the deme reached 500 updates. To promote movement, individual organisms were rewarded with 10 units of energy for each of their first 5 moves. Targets comprising one or more cells were defined and persisted for the duration of the run. Because organisms were injected with random facings into random cells in a toroidal environment, no recurring path could be exploited between an organism's injection cell and the target.

Evolution of Effective Search Strategies. In our initial set of runs, the deme replication predicate required 8 of the 10 injected organisms to locate a 2x2-cell target in a 40x40 deme. A common strategy evolved that did not include the use of pheromones, but was nonetheless highly successful at completing the task by efficiently searching the environment. An example of these search patterns is shown in Figure 4. This strategy consisted of a number of moves in one direction, a rotation, a number of additional moves, and then a rotation back to the original heading. The main axis of movement was almost always along a diagonal. The periodic "kink" in the pattern allows the organism to break from the current row, column, or diagonal and explore new areas, eventually finding a path leading to the target. Because these evolved organisms were so successful at searching for the target, there was little pressure for them to use pheromones to help each other complete the task. Successful strategies emerged after 10,000-20,000 updates, allowing demes to replicate more frequently than the 500-update age limit.

Testing a Hand-Built Organism. To determine whether pheromone use would improve upon the per-

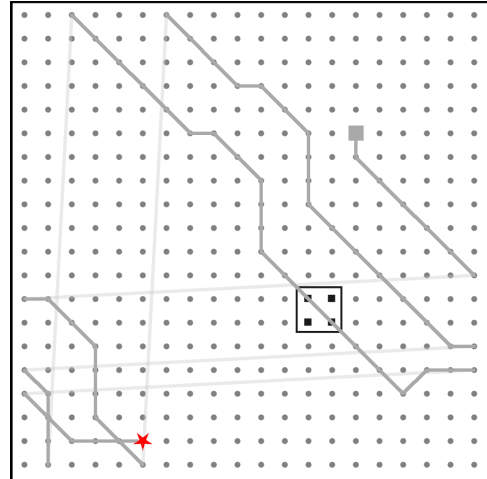


Fig. 4. Typical search pattern for an evolved organism. The star indicates the organism's injection cell, the square represents its current position, and lighter lines indicate movements that wrap around the torus.

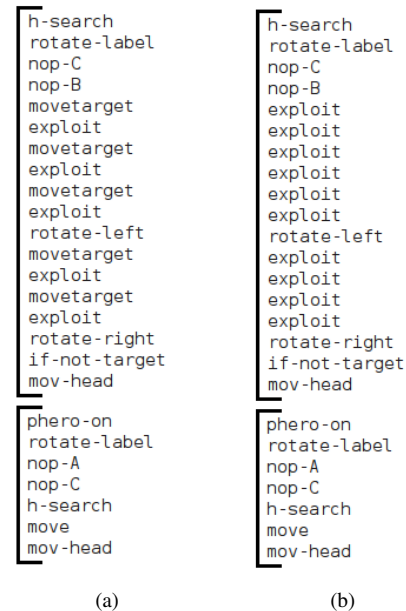


Fig. 5. Genomes of hand-built organisms. The upper portion searches the environment for target cells, continuing along a "kinked" diagonal path until it reaches a target cell or a pheromone trail laid by another organism. In the latter case, it follows the pheromone trail to the target. In the lower portion, after the target has been reached, the organism turns on pheromone dropping, rotates to face a perpendicular diagonal, and moves forward repeatedly. In algorithm (a), the search consists of alternating `exploit` and `movetarget` instructions, while search in algorithm (b) consists entirely of `exploit` instructions.

formance seen in these runs, an organism was created by hand that followed the evolved search strategy, but also used pheromones to mark a trail to the target. The genome of this organism is shown in Figure 5(a). This organism has two states, each of which is contained in a loop. If it has not yet found the target, it will continue to make 6 movements along the diagonal with an occasional kink consisting of 4 movements, return to the diagonal,




Fig. 6. Movements and pheromone dropping of the second hand-built organism. Stars indicate organisms' injection cells, squares represent the current position of organisms, thin colored lines represent movements without pheromone dropping, and bold gray lines represent movements with pheromone dropping. Individual cells are marked with circles if they have zero pheromone or squares if they have pheromone.

and repeat this process until it reaches a target cell. If the organism encounters a pheromone trail, `exploit` will pull the organism towards the strongest source. Once the organism reaches a target cell, it proceeds to the lower section of its genome, which turns on pheromone dropping, rotates the organism to face the diagonal perpendicular to the search direction, and continually moves the organism forward, wrapping around the torus. The goal of this behavior was to create a straight pheromone trail that leads directly to the target.

Runs seeded with this hand-built organism were successful at finding the target, and the use of pheromones persisted in the population. However, many organisms had difficulty following the pheromone due to the presence of parallel pheromone trails.

Improved Hand-Built Organism. To address these deficiencies, a second hand-built organism was constructed. This organism, shown in Figure 5(b), emphasizes `exploit`-based movements. Target cells were artificially marked with pheromone. This genome was successful at performing the task, maintained the use of pheromone in its genome, and did not encounter problems with moving through or breaking away from a pheromone trail. Figures 6 and 7 show an example of the movements and pheromone levels, respectively, of one deme containing organisms with this genome.

Solutions Evolved from Scratch. Several runs were conducted to test whether evolution would remove or improve upon the pheromone use in the updated hand-built organism. For comparison, additional runs were performed in identical environments that started from scratch (i.e., a genome containing only `nop-C` instructions, which have no effect on the organism or environ-

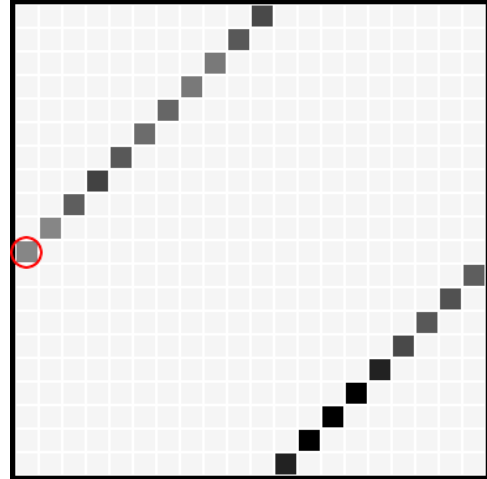


Fig. 7. Distribution of dropped pheromone levels by hand-built organisms. Darker cells indicate higher levels of pheromone. The target cell is located in the center of the left side, indicated with a red circle.




Fig. 8. Mean duration in updates (\pm SD) required for organisms in a deme to reach the target. Organisms either evolved from scratch or from the updated hand-built organism. The first set of data is for organisms that did not have access to the pheromone resource or instructions.

ment). The environment for these runs was a 20x20 torus containing a 1x1 target, making the task slightly more difficult, and allowing runs to be completed more quickly. As seen in earlier runs, demes began exhibiting successful strategies within 10,000-20,000 updates.

To compare the effectiveness of the solutions evolved over 50,000 updates, the resulting dominant organisms were used to seed 100 demes with 20 organisms each, and the number of updates and total number of moves required for 16 unique organisms to reach the target were recorded. In these tests, mutations and replication were turned off, so demes either accomplished the task or died.

Figures 8 and 9 show the mean number of updates and organism movements required, respectively, for the task to be satisfied. We observe that each run produced organisms




Fig. 9. Mean total movements (\pm SD) required for organisms in a deme to reach the target. Organisms either evolved from scratch or from the updated hand-built organism. The first set of data is for organisms that did not have access to the pheromone resource or instructions.

that were able to complete the task in shorter average number of updates and with fewer total movements when using pheromones than runs in which organisms did not have access to the pheromone resource. Moreover, organisms that evolved from scratch were able to complete the task in a shorter time and with fewer movements, on average, than those that evolved from the hand-built solution. Common among these successful organisms was the use of pheromones around the target area rather than a single trail leading to the target. Effectively, pheromone levels built up around the target and radiated outward, as shown in Figure 10. This behavior has the effect of enlarging the target, so that once an organism senses pheromone, it can simply follow it into the target. The dominant organisms all displayed patterns of movements that contained several movements in one direction followed by a rotation, almost always using `explore`. Because of the random rotation included in `explore`, the trajectories of these organisms were not as predictable as those seen previously, but they were, nonetheless, successful at reaching the target cell. It is possible that organisms abandoned the 6-diagonal, 4-kink movements of the hand-built organism, because it was evolved for a 2×2 target, while the target in these runs was 1×1 . We note that the more random search pattern is likely to be more robust to differing environments.

In a small number of runs, organisms turned on pheromone dropping once they encountered a cell with pheromone as opposed to seeing the target. This behavior enables paths to be re-enforced more quickly, but it also has the downside of potentially re-enforcing inefficient trails. It is suspected that this response emerged as a result of the target being painted with pheromone.




Fig. 10. Distribution of pheromone resource around a 1×1 target cell, located in the left center, after 8 updates. The cell with the highest amount of pheromone is the target.

VI. CONCLUSIONS AND FUTURE WORK

This work explored the evolution of pheromone-based communication in groups of digital organisms. Given the task of getting members of a population to a target, initial populations developed an efficient search strategy that greatly mitigated the benefits of pheromones. Organisms in an environment with a more elusive target, however, evolved a collaborative strategy that effectively shrank the search space by dropping pheromones in a pattern that radiated outwards from the target. Group members were able to follow these patterns into the target. These evolved solutions proved to be more robust than hand-built solutions.

This research can be extended in several ways. First, pheromone evaporation in runs has thus far been limited to small rates, typically 0%, and 10%, which yielded similar strategies. Runs with higher evaporation rates might lead to more liberal use of pheromones, as mistakes will more quickly disappear from the environment. This flexibility could result in the emergence of exploratory patterns, aid in deciding among multiple targets, or promote a division of labor or pro-active and re-active foraging [36]. Second, in the natural world, organisms use pheromones to navigate both to and from a target. Such behavior would not be possible with the `exploit` instruction, which always follows the pheromone gradient upwards. Additional capabilities, such as the ability to know whether or not the organism is currently carrying “food,” could promote these behaviors. Alternately, the evolution of a memory, landmark recognition, and other forms of navigation may be critical to this task.

FURTHER INFORMATION

Additional materials related to digital evolution and Avida are available at <http://devolab.msu.edu>, and information regarding the evolution adaptive and cooperative behaviors can be found at <http://www.cse.msu.edu/thinktank>.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the contributions of Dave Knoester and Charles Ofria. This work was supported by NSF Grants CCF-0750787, CNS-0751155, CCF-0820220, by U.S. Army Grant W911NF-08-1-0495, and by a Quality Fund Grant from Michigan State University.

REFERENCES

- [1] B. Bassler, "Small talk cell-to-cell communication in bacteria," *Cell*, Jan 2002.
- [2] F. Matsumura, H. Coppel, and A. Tai, "Isolation and identification of termite trail-following pheromone," *Nature*, vol. 219, pp. 963–964, August 1969.
- [3] A. Granero, J. Sanz, F. Gonzalez, J. Vidal, A. Dornhaus, J. Ghani, A. Serrano, and L. Chittka, "Chemical compounds of the foraging recruitment pheromone in bumblebees," *Naturwissenschaften*, vol. 92, pp. 371–374, Jan 2005.
- [4] R. Zimmer, "Chemical signaling processes in the marine environment," *The Biological Bulletin*, vol. 198, no. 2, pp. 168–187, 2000.
- [5] L. Buck, "The Molecular Architecture of Odor and Pheromone Sensing in Mammals," *Cell*, vol. 100, no. 6, pp. 611–618, 2000.
- [6] P. Karlson and M. Luscher, "Pheromones - new term for a class of biologically active substances," *Nature*, vol. 183, pp. 55–56, Jan 1959.
- [7] B. Hölldobler and E. Wilson, *The Ants*. Harvard University Press, 1990.
- [8] C. Ofria and C. Wilke, "Avida: A software platform for research in computational evolutionary biology," *Journal of Artificial Life*, vol. 10, pp. 191–229, Jan 2004.
- [9] A. Traulsen and M. A. Nowak, "Evolution of cooperation by multilevel selection," *Proceedings of the National Academy of Sciences*, vol. 103, pp. 10952–10955, Jan 2006.
- [10] M. A. Nowak, "Five rules for the evolution of cooperation," *Science*, vol. 314, pp. 1560–1563, Jan 2006.
- [11] E. O. Wilson and B. Hölldobler, "Eusociality: Origin and consequences," *Proceedings of the National Academy of Sciences*, vol. 102, pp. 13367–13371, Jan 2005.
- [12] M. Symonds and M. Elgar, "The evolution of pheromone diversity," *Trends in Ecology & Evolution*, Jan 2008.
- [13] W. Roelofs and A. Rooney, "Molecular genetics and evolution of pheromone biosynthesis in Lepidoptera," *Proceedings of the National Academy of Sciences*, vol. 100, pp. 9179–9184, Jan 2003.
- [14] W. Roelofs, W. Liu, G. Hao, H. Jiao, and A. Rooney, "Evolution of moth sex pheromones via ancestral genes," *Proceedings of the National Academy of Sciences*, vol. 99, pp. 13621–13626, Jan 2002.
- [15] D. Chivers, G. Brown, and R. J. Smith, "The evolution of chemical alarm signals: Attracting predators benefits alarm signal senders," *American Naturalist*, vol. 148, pp. 649–659, Jan 1996.
- [16] R. Emes, S. Beatson, C. Ponting, and L. Goodstadt, "Evolution and comparative genomics of odorant-and pheromone-associated genes in rodents," *Genome Research*, vol. 14, pp. 591–602, Jan 2004.
- [17] H. Eisthen, "Why are olfactory systems of different animals so similar?," *Brain, Behavior and Evolution*, vol. 59, pp. 273–293, Jan 2002.
- [18] E. Mallon, S. Pratt, and N. Franks, "Individual and collective decision-making during nest site selection by the ant *Leptothorax albigipennis*," *Behavioral Ecology and Sociobiology*, vol. 50, pp. 352–359, Jan 2001.
- [19] D. Sumpter and M. Beekman, "From nonlinearity to optimality: Pheromone trail foraging by ants," *Animal Behaviour*, vol. 66, pp. 273–280, Jan 2003.
- [20] R. Beckers, O. Holland, and J. Deneubourg, "From local actions to global tasks: Stigmergy and collective robotics," in *Proceedings of Artificial Life IV*, Jan 1994.
- [21] M. Dorigo and G. Di Caro, "Ant colony optimization: A new metaheuristic," in *Proceedings of the 1999 Congress on Evolutionary Computation*, 1999, vol. 2, 1999.
- [22] G. D. Caro, F. Ducatelle, and L. M. Gambardella, "AnthocNet: An ant-based hybrid routing algorithm for mobile ad hoc networks," *Lecture Notes in Computer Science*, Jan 2004.
- [23] R. J. Collins and D. R. Jefferson, "Antfarm: Towards simulated evolution," in *Artificial Life II*, pp. 579–601, Jan 1991.
- [24] L. Panait and S. Luke, "Learning ant foraging behaviors," in *Proceedings of the Ninth International Conference on Artificial Life*, Jan 2004.
- [25] J. A. Sauter, R. Matthews, H. V. D. Parunak, and S. Brueckner, "Evolving adaptive pheromone path planning mechanisms," in *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems*, Jan 2002.
- [26] D. Floreano, S. Mitri, S. Magnenat, and L. Keller, "Evolutionary conditions for the emergence of communication in robots," *Current Biology*, vol. 17, pp. 514–519, 2007.
- [27] R. Lenski, C. Ofria, R. T. Pennock, and C. Adami, "The evolutionary origin of complex features," *Nature*, Jan 2003.
- [28] B. E. Beckmann, P. K. McKinley, and C. Ofria, "Evolution of an adaptive sleep response in digital organisms," *Lecture Notes in Computer Science*, vol. 4648, pp. 233–242, 2007.
- [29] D. Knoester, P. K. McKinley, B. E. Beckmann, and C. Ofria, "Directed evolution of communication and cooperation in digital organisms," *Lecture Notes in Computer Science*, Jan 2007.
- [30] P. K. McKinley, B. Cheng, C. Ofria, D. Knoester, B. E. Beckmann, and H. Goldsby, "Harnessing digital evolution," *IEEE Computer*, vol. 41, Jan 2008.
- [31] D. B. Knoester and P. K. McKinley, "Cooperative network construction using digital germlines," in *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, July 2008.
- [32] L. Grabowski, W. Elsberry, C. Ofria, and R. T. Pennock, "On the evolution of motility and intelligent tactic response," in *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, July 2008.
- [33] B. E. Beckmann and P. K. McKinley, "Evolution of adaptive population control in multi-agent systems," in *Proceedings of the Second IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO)*, (Venice, Italy), October 2008.
- [34] D. Wilson, "Altruism And Organism: Disentangling The Themes Of Multilevel Selection Theory," *The American Naturalist*, vol. 150, no. S1, pp. 122–134, 1997.
- [35] B. D. Connelly, P. K. McKinley, and B. E. Beckmann, "Evolving cooperative pheromone usage in digital organisms," Tech. Rep. MSU-CSE-08-32, Computer Science and Engineering, Michigan State University, East Lansing, Michigan, November 2008.
- [36] F. Dechaume-Moncharmont, A. Dornhaus, A. Houston, J. McNamara, E. Collins, and N. Franks, "The hidden cost of information in collective foraging," *Proceedings of the Royal Society B: Biological Sciences*, vol. 272, no. 1573, pp. 1689–1695, 2005.