

Experimenting with a Real-Size Man-Hill to Optimize Pedagogical Paths

G. Valigiani
LIL Department
University of Calais
valigian@lil.univ-littoral.fr

Y. Jamont
Paraschool company
yannick.jamont@paraschool.com

C. Bourgeois
Republique
LERSIA Department
University of Bourgogne
bourgeois.republique@free.fr

R. Biojout
Paraschool company
raphael.biojout@paraschool.com

E. Lutton
Complex Team
INRIA Rocquencourt
evelyne.lutton@inria.fr

P. Collet
LIL Department
University of Calais
collet@lil.univ-littoral.fr

ABSTRACT

This paper describes experiments aimed at adapting Ant Colony Optimization (ACO) techniques to an e-learning environment, thanks to the fact that the available on-line material can be organized in a graph by means of hyperlinks between educational topics. The structure of this graph is to be optimized in order to facilitate the learning process for students.

ACO is based on an ant-hill metaphor. In this case, however, the agents that move on the graph are students who unconsciously leave pheromones in the environment depending on their success or failure. In the paper, the whole process is therefore referred to as a “man-hill.”

Compared to the [13, 14] papers that were providing guidelines for this problem, real-size tests have been performed, showing that man-hills behave differently from ant-hills. The notion of pheromone *erosion* (rather than evaporation) is introduced.

1. INTRODUCTION

Back in 2002, Paraschool, the French leading e-learning company was looking for a system that could enhance site navigation by making it intelligent and adaptive to the user.

Since their software is based on a graph traversed by students (where pedagogical items are nodes and hypertext links are arcs), ACO techniques [6, 1, 2] can apply and show interesting properties: adaptability and robustness.

ACO (developed after the observation of ant-hills [8, 4]) uses virtual ants to find minimal paths in a graph. In the Paraschool system, more than 100,000 students explore the graph. The very large number of students triggered the idea to apply virtual ant-hill techniques using real students rather than virtual ants, with the aim of optimizing pedagogical paths traversing a set of educational topics.

Two papers [13, 14] have already described the first steps

toward the use of ACO-like algorithms with humans in the e-learning domain. Real-size experimentations have shown that ant-hill optimization techniques developed in ACO do not directly apply because students do not behave like artificial ants.

This paper uses real-size experiments to present the notion of what is termed “man-hill optimization” and its specificities w.r.t. ant-hill optimization. Results are studied and a conclusion discusses the man-hill paradigm and its future developments.

2. IMPLEMENTATION OF THE PARASCHOOL “MAN-HILL”

The Paraschool e-learning software is used in French schools or by individual students at home over the Internet. Connected students have access to thousands of pedagogic items (know-hows, lessons, drills) that were originally deterministically related by hypertext links.

The aim of this work is twofold: find the best succession of items so as to optimize learning and insert some intelligence into the system so that different students have a different view of the Paraschool software.

Rather than using artificial ants like in ACO, the very large number of users makes it realistic to use them directly to release pheromones on the graph, depending on how they validated an item (success or failure).

2.1 Arc selection

In the original Paraschool software, when an item was validated, students were deterministically presented another item until the lesson was completed. The graph structure was in fact a collection of linked lists.

Using an ACO-like algorithm means that lessons, know-hows and drills need to be structured into a real graph, giving the student several potential choices after an item is validated.

Arcs interconnecting items bear values that are used to compute a “fitness” function that evaluates how nice it would be for a particular student to follow a particular arc. This values can be seen as probabilities for a student to follow an arc.

A selection method borrowed from genetic algorithms¹ is then used to choose which arc will be followed when the

¹Tournament, stochastic tournament, ranking, ... [10, 13, 14].

student clicks on the **NEXT** button.

2.2 Biasing the graph w_p

In the Paraschool man-hill, even though the emergence of good paths is what is sought, one must also take pedagogy into account. The arcs of the Paraschool graph are therefore initialized by the pedagogic team with relative values they think right, meaning that arcs are not born equal. When a teacher creates a new theme, he is asked to relate different topics with valued arcs representing the probability he thinks right for a student to follow an arc.

Relative Pedagogic Weight w_p : In order to avoid discrepancy between the way different teachers rate the arcs they create, suggested weights are converted to relative values. If w_1, \dots, w_n are the given weights of n outgoing arcs, the relative pedagogic weight of arc j is $w_j / \sqrt{\sum_{i=1}^n w_i^2}$.

To simplify their task even more, teachers are asked to choose values within $[0, 20]$ to represent the permanent bias they want to initialize the graph with.

2.3 Paraschool man-hill pheromones

The first idea in order to get good pedagogic paths emerge from the graph has been to follow a standard ant-hill ACO paradigm, but with several types of pheromones.

2.3.1 Global pheromones

Ant-like positive (φ^+) and negative (φ^-) pheromones are released on arcs, depending whether students succeeded or failed on the item pointed to by the arc.

These student pheromones are added to the relative pedagogic weight (w_p) in order to determine the fitness of an arc.

2.3.1.1 Arc creation:

If a student jumps from one item to another without using an existing arc (using the graph browser, for instance), a new arc with a default pheromone value of 1 is automatically created.

Therefore, if many students decide to jump to another topic rather than follow the teacher's advice, the initial w_p given by the teacher will lose its efficiency. The pheromone on the created arc fades away so that rebellious arcs do not accumulate in time.

The construction of new arcs is important, because it might indicate a problem with the exercise from which they shoot: The creation of many arcs from a physics exercise toward a maths trigonometry section may indicate that students have encountered trigonometry problems while solving this exercise. The pedagogical team may use this information to take correcting actions.

2.3.2 Individual pheromones

Paraschool also wanted their software to adapt to each student. It appeared that this request was not only desirable but also necessary: A student may not like being proposed to solve the same exercise several times, for instance.

Thus special kinds of "individual" pheromones that has been and should be implemented only belong to one student. They are pheromones because they are left on arcs and they evaporate. However, the stigmergic information

they implement is only used by the student who released them.

The first implemented individual pheromone is the History Weight (φ_h). This pheromone is used as a multiplicative factor over the $w_p + \varphi^+ + \varphi^-$ sum. It is worth noting that for a multiplicative pheromone, evaporation means going back to 1 (neutral element for multiplication) *and not back to 0* (neutral element for addition) as for additive pheromones.

When a student validates an item, a φ_h pheromone is released with a value of 0.1, for instance. This divides by 10 the probability for a student to be proposed an item (s)he recently validated. However, φ_h goes back to 1 as it evaporates, so that the probability for this item to appear again increases with time, which may not be a bad idea since student memory is also known to fade away with time.

Other individual pheromones can be implemented (an agenda pheromone to force students to validate items chosen by their teacher, ...).

2.4 Arc fitness function

The arc fitness function makes use of all the available information, along with tuning parameters (w_i) that can be changed by the Paraschool pedagogic team:

$$arc_fitness = \varphi_h \cdot (w_1 \cdot w_p + w_2 \cdot \varphi^+ - w_3 \cdot \varphi^-)$$


where φ^{+-} are the pheromones, w_p the pedagogical weight, φ_h the history pheromone and w_i the tuning parameters.

3. EXPERIMENTING WITH THE PARASCHOOL MAN-HILL

3.1 A simple test case

In order to adjust the system, the common experiment used to show the adaptability of natural ants to find a shortest path [4] was adapted to the pedagogical environment (cf. fig. ??). A simple graph contains four items: a starting item S , two intermediate items I_b/I_g and an ending item E .

Figure 1: Verifying the possibility of arc fitness inversion on a simple test case.



Let us suppose that Paraschool teachers feel that the best way to validate item E when coming from item S is to go through item I_b rather than going through item I_g . They bias the graph by giving relative weights of 15/20 to the top arcs, and 1/20 to the bottom arcs.

If the teachers happen to be wrong, this is typically what the system should be able to correct: on the graph of fig. ?? one observes a success rate of 90% on items I_b and I_g when coming from item S , and a success rate of 90% or 10% on item E depending whether students have gone through item I_g or I_b respectively.

This simple experiment has been used to set the first values of the man-hill global parameters. Weights w_1 , w_2 and w_3 have been adjusted respectively to 1, 1 and 3 so that the fitness value of the path established by teachers could be overcome by the student path (fitness inversion). Other

tests have led to setting an evaporation rate of 0.999 per day and the release of pheromones per student to 0.1 (cf. [14]).

3.2 Analysis of on-line experiments

The [14] paper ended by saying that on-line experiments would be starting soon. During twelve months, the algorithm was implemented and tested on the real Paraschool graph with the aim of gathering as much data as possible.

By default, all arcs had been set to 1 so that the system would appear to work “normally,” while the pedagogic team started modifying the graphs theme by theme, in maths. Around 4 000 arcs were manually updated or created. During this period, the algorithm was started and set to run in “silent mode,” i.e. only gathering pheromones and allowing the creation of new “student arcs,” without offering to the user any man-hill suggestions: it was important to make sure that the real-size implementation of this system would not lead to unacceptable CPU or memory load, or database size.

During the longest test period, students had logged on the Paraschool software more than 70.000 times. More than 30.000 student arcs had been created. Half of them had been followed only once, while the other half had been followed 4 times in average.

In average, the success rate per item was 64%. Another interesting thing is that even though after 12 months, the average pheromone value on each arc was 0.37, one arc presented a huge quantity of pheromone of 27.

In fact several abnormal values have been found during the test period that were tracked down to problems in the Paraschool graph. This is a very positive result on the use of the man-hill concept as an automatic audit tool.

3.3 Study on the evaporation rate

The collected data allowed to playback student visits on the Paraschool graph for a 12 months period, which is very interesting, since this allowed to test several options. Originally, evaporation was implemented as multiplying the value of a pheromone by 0.999 per day. This value had been determined on the simple test case (cf. 3.1) so that the fitness inversion could occur.

On the real system, inversion occurs rapidly, but evaporation is very slow (cf. fig. ??a). The system diverges as pheromones accumulate a lot, leading to a state where it becomes impossible for another inversion to occur again if conditions change. The adaptive characteristics of ACO is lost.

Fig. ??b shows a simulation with a quicker evaporation rate of 0.974. Pheromones do not accumulate as fast as with 0.999, even though fitness on the bottom arc still gets very high in October. The system now retains adaptive characteristics but one sees that during the summer holiday, pheromones evaporate quickly, meaning that although test conditions have not changed, the wrong path set up by the teacher becomes prominent again.

This leads to a big problem as none of the simulations are satisfactory: the first one is too stable for the system to be adaptive while the second one forgets obtained results if not used frequently over a period of time. Such a phenomenon is not frequently encountered in artificial ant-hills which are not linked to real-time constraints, like seasonal activities.

3.4 Erosion: a new “pheromone evaporation” concept better suited to man-hills

An adaptive evaporation rate depending on theme activity could be imagined, but some themes are interconnected, making it difficult to arbitrarily partition the Paraschool graph into activity-disconnected themes.

Another possible approach is to implement an evaporation per arc, which could be seen as an erosion mechanism: when following an arc, an ant erodes the pheromones borne by the arc, with the interesting result that erosion only occurs on used arcs. One large drawback of this method is that if an arc is wrongfully handicapped because it was first followed by bad students, a vicious circle may appear: the current arc will be rarely chosen and will not be given much opportunity to change its bad fitness value.

A less radical approach is to implement erosion on a per-item basis: when following an arc, erosion occurs on *all* arcs shooting from the item. Pheromone erosion of a non-chosen arc is the price to pay for not having been elected, while the followed arc will receive whatever pheromones are earned by the validation of the pointed item.

This extended erosion method presents both the advantage of not losing information during long periods of inactivity (no time-based evaporation) while still preserving the essential pheromone evaporation process on non-used arcs.

3.5 What should be the goal of the Paraschool man-hill ?

After several months of silent-mode testing, Paraschool has been satisfied enough with the robustness of the system to allow man-hill suggestions to be proposed to students.

The original aim of the system was to maximize success (arc fitness calculation favored arcs leading to successfully validated items) (cf. [13, 14]).

Results have been quite interesting as they showed that this original aim was a bit naive.

A precise analysis of theme 377 (Mathematics, 6th Grade: Decimal numbers on which man-hill suggestions were activated) revealed that this is exactly what was happening: the main paths containing the arcs with the greatest pedagogical weights were quickly neglected in favor of a path which quickly brought the students successfully to the end of the theme. . .

The Paraschool man-hill had found an optimal path toward “success” where all items of the path were easy to validate.

Unfortunately, the pedagogical team pointed out that this was not exactly what teaching was about, so the fitness function had to be refined to try to maximize learning rather than to maximize success.

3.6 A new fitness function to maximize learning

Originally, maximizing success had been achieved by adding $w_2.\varphi^+$ (success pheromone) and subtracting $w_3.\varphi^-$ (failure pheromone) in the arc fitness function. In order to maximize *learning*, the Paraschool pedagogic team asked that the system should try to find paths where items would be moderately hard to validate.

The suggestion was then to favor arcs leading to items on which average students would have 60% chances of succeeding and 40% chances of failing. The 60/40 ratio was chosen so that students would succeed slightly more than

they would fail, so that they are not discouraged by failing too often.

The $w_2 \cdot \varphi^+ - w_3 \cdot \varphi^-$ term was therefore replaced by $w_2 \cdot (\frac{60}{40} - \frac{\varphi^+}{\varphi^-})$, which represents the inadequacy to the desired success rate.

Then, another refinement was suggested: when the amount of pheromones on the arcs is not large enough to bear significant enough information, Paraschool asked that the relative pedagogic weight set up by the teachers be used regardless of student pheromones.

3.7 Current fitness function

The new arc fitness function reduces the pedagogical weight proportionally to the amount of pheromone borne by the arc. When enough pheromone is present, maximizing the 60/40 success rate becomes the only important thing.

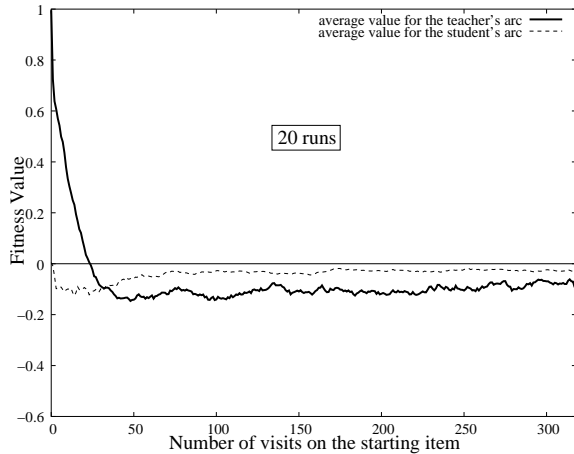
The expression of the new fitness function is now defined by $\varphi_h \cdot w_p \cdot \max(1 - w_1 \cdot (\varphi^+ + \varphi^-), 0) - \min(w_2 \cdot (\frac{60}{40} - \frac{\varphi^+}{\varphi^-}), w_3 \cdot (\frac{\varphi^+}{\varphi^-} - \frac{60}{40}))$

where $\varphi^{+/-}$ are the pheromones, w_p the pedagogical weight, φ_h the history pheromone and w_i the tuning parameters. w_1 allows to tune the amount of pheromones above which w_p is not used any more (its value is 0.85). w_2 and w_3 have been dissociated in order to allow having different slopes depending on whether the arc is bad because it leads to an item too difficult or too easy to validate. $w_2 = 0.4$ and $w_3 = 1$ because it is considered to be better for an arc to lead to an easy item rather than to a difficult one.

4. OFF-LINE EXPERIMENTS WITH THE NEW FITNESS FUNCTION AND PHEROMONE EROSION

The small test case used in section 3.3 has been modified to suit the new objectives. Fig. 2 shows on an average of 20 runs the fitness of teacher arc (bold) and student arc (dashed), when the teacher advises the good arc (fig. 2.a) or the wrong arc (fig. 2.b) with a pedagogical weight of 15.

Figure 2: Behavior of the system on a small test case during 20 runs. On the left figure, the teacher (bold curve) has advised the good arc, while on the right figure, he has advised the wrong one.



Since erosion is implemented, the X axis does not represent time anymore, but a number of visits (318 in this example).

As expected, the bold fitness value decreases constantly during the first 40 visits until the pheromone information is considered as significant.

Then, depending whether the teacher advised the right or wrong arc, both arcs stay away one from each other, with an inversion when the wrong arc was advised.

Difference between both arcs is smaller on fig. 2.b. because when the man-hill determined that following the teacher's advice lead to success rates incompatible with Paraschool's policy (60/40), the arc is less used, and the amount of pheromone it bears erodes without being reinforced. When it gets below the w_1 threshold, pedagogic weight becomes prevalent again and it takes some time for students to discover again that this arc was not a good one to follow.

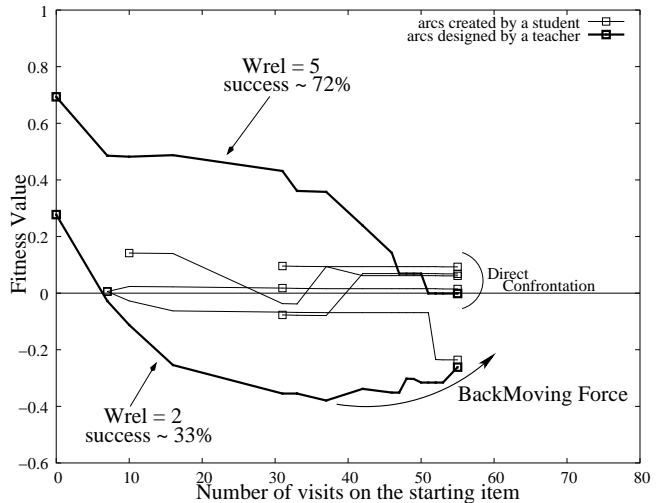
All in all, in average, the system behaves as expected (inversion happens). Moreover, it can be seen as a positive thing that even though the teacher had advised a wrong arc, the system tends to keep proposing it more often than if it had not been advised.

5. ON-LINE OBSERVATIONS

The new fitness function and erosion method have been set on-line for one month on the Paraschool software.

Fig.3 shows the fitness value evolution of outgoing arcs where two teacher arcs correspond to the two different cases.

Figure 3: The top bold curve belongs to an arc that leads to an item with a success rate of around 70%, while the bottom one leads to an item with a success rate of around 33% only.



The bold curves represent teacher arcs and the thin curves represent arcs created by students. The upper teacher arc has been advised rightly. Its pedagogical weight has been set to 5 and its success rate is for the moment around 70%, which approximately corresponds to the 60/40 ratio established by Paraschool.

On the contrary, the second teacher arc leads to an item with a success rate of only 33% and its pedagogical weight is 2.

An inversion occurs rapidly on the second arc, while on the first, pedagogic influence decreases with the number of visits, up to the point where its fitness struggles with those of the arcs found by students.

On the last 15 visits, the bad teacher arc is not used very often, and the pheromones it bears are eroded by students choosing other arcs. Pedagogic weight becomes more prevalent up to the point where students will follow it again and find again that this is not a good arc to follow.

If this happens repeatedly, thanks to the monitoring system, the Paraschool pedagogic team should detect the wrong arc and should either modify the item to which the arc leads, or should advise another item. The system should be seen as an auditing tool that helps teachers build relevant pedagogical paths.

Compared to the toy problem, the system did not react as quickly as expected. This comes from the fact that students create several other arcs (not an existing option in the test case) which attract pheromones that are not released on the teacher arc, therefore delaying confrontation.

6. CONCLUSIONS AND DEVELOPMENTS

Developing an ant colony optimization technique using human students on the Paraschool graph has led to the conclusion that artificial ant-hills and natural man-hills do not behave exactly in the same way:

- Usually, artificial ants are programmed so as to solve the problem at hand. On the contrary, human students do not follow a particular algorithm. They go wherever they like and can only be suggested to follow a particular path (which might be closer to real ants behavior).
- Artificial ants are permanently active on the entire environment, to the contrary of students, who go periodically on holiday, and are studying different topics along the year. Standard artificial ant-hills “time”-based evaporation techniques lead to information loss when used on real man-hills, because students are not equally active all year long, on all parts of the graph. A new erosion concept has been developed that seems to be better adapted to man-hills and that could also be used on artificial ant-hill problems that are faced with uneven activity.
- Constraints will apply on human students that simply do not make sense with artificial ants. Artificial ants, for instance, will never complain of going several times through the same arcs. Human students, on the contrary, will immediately recognize a drill that is proposed for a second time and may get fed up with the system if this happens too often.
- Natural ant-colonies and other social insects implement some kind of altruistic behavior, where the future of the colony is more important than a single individual. Individual students may not like to get lost into some remote branch of the software just for the sake of exploring the environment.

Because of these differences, standard ACO algorithms do not work straight out of the box. In fact, man-hill optimization is different from ant colony optimization, and it seems a very interesting research topic indeed: if social insects are able to find optimal paths or structures thanks to emergence and stigmergy, the same process clearly takes place in human societies.

As with insects, emergence processes are still mostly unconscious in human societies. Many structures have emerged over the Internet without any clear design. Studying the behavior of what we call “man-hills” could allow to harness this power to optimize desired features. Many commercial

software are now available over the net that have thousands of users browsing their sites per day. Such companies could use the emerging capabilities of “man-hills.”

With Paraschool, the concept has already shown its power at least as an automatic auditing tool on huge graphs that are too large to be maintained by hand (pheromones rates rise to abnormal levels on arcs leading to erroneous items). It can also help site designers with suggestions automatically emerging from users browsing the site.

A strong and stable basis now seems to be available for the Paraschool software, where new concepts such as pheromone erosion and multiplicative pheromones eroding back to 1 have been introduced. The man-hill experiment will now scale to the whole graph containing thousands of nodes. The next step will now be to study the “optimal learning paths” found by students, and try to understand the complex behavior of man-hills.

7. REFERENCES

- [1] E. Bonabeau, M. Dorigo, G. Theraulaz, *Swarm Intelligence : From natural to Artificial systems*, Oxford University Press 1999, ISBN 0-19-513159-2.
- [2] E. Bonabeau, M. Dorigo and G. Theraulaz, *Inspiration for optimization from social insect behavior*, in *Nature*, vol. 406 pp 39–42, 2000.
- [3] A. Colomi, M. Dorigo and V. Maniezzo, *Distributed optimization by ant-colonies*, in proceedings of European Conference on Artificial Life, Cambridge, MIT Press, pp 134–142, 1991.
- [4] J.L. Deneubourg, S. Aron, S. Goss and J.M. Pasteels, *The self-organizing exploratory pattern of the argentine ant*, in *Journal of Insect Behavior*, vol. 3 pp 159–168, 1990.
- [5] M. Dorigo, E. Bonabeau and G. Theraulaz, *Ant algorithms and stigmergy*, in *Future Generation Computer Systems*, vol. 16 pp 851–871, 2000.
- [6] M. Dorigo and G. Di Caro, *The ant colony optimization metaheuristic*, in *New ideas in optimization*, D. Corne, M. Dorigo and F. Glover (Eds), McGraw-Hill, pp 11–32, 1997.
- [7] M. Dorigo, *Optimization, learning and natural algorithms*, PhD Thesis, politecnico di Milano, 1992.
- [8] J.L. Deneubourg, J.M. Pasteels and J.C. Verhaeghe, *Probabilistic behavior in ants : a strategy of errors?*, in *Theoretical Biology*, vol. 105 pp 259–271, 1983.
- [9] D.E. Goldberg and K. Deb, *A comparative analysis of selection schemes used in genetic algorithms*, in G. Rawlins, editor, *Foundations of Genetic Algorithms*, vol. 1 pp 69–93, 1991.
- [10] D.E. Goldberg, *Genetic algorithms in search, optimization and machine learning*, Addison-Wesley Publishing Company Inc., Reading, MA, 1989.
- [11] N. Labroche, N. Monmarch and G. Venturini, *A new clustering algorithm based on the chemical recognition system of ants*, in proceedings of the European Conference on Artificial Intelligence, IOS Press, pp 345–349, 2002.
- [12] M. Resnick, *Turtles, termites and traffic jams : Explorations in massively parallel microworlds*, Complex adaptive systems series MIT Press, 1994.
- [13] Y. Semet, Y. Jamont, R. Biojout, E. Lutton and P. Collet, *Artificial Ant Colony and E-Learning : An optimization of pedagogical paths*, HCI 2003.
- [14] Y. Semet, E. Lutton and P. Collet, *Ant Colony Optimization for E-Learning : Observing the emergence of pedagogic suggestions*, SIS 2003.
- [15] T. Sttzle, M. Dorigo, *ACO Algorithms for the travelling salesman problem*, in proceedings of the EUROGEN conference, M. Maleka, K. Miettinen, P. Neittaanmaki, J. Periaux (Eds), John Wiley & Sons, pp 163–183, 1999.