Adaptation of the ACO heuristic for sequencing learning activities

Sergio Gutiérrez[1], Grégory Valigiani[2], Pierre Collet[2] and Carlos Delgado Kloos[1]

[1]University Carlos III of Madrid (Spain)
[2]Université du Littoral (France)
[1]{sergut,cdk}@it.uc3m.es, [2]{valigian,collet}@lil.univ-littoral.fr

Abstract. This paper describes an initiative aimed at adapting swarm intelligence techniques (in particular, Ant Colony Optimization) to an e-learning environment, thanks to the fact that the available online material can be organized in a graph by means of hyperlinks of educational topics. In this case, 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 referred as man-hill, as opposed to the ant-hill metaphor of ACO. The paper presents the system and shows the experimental results obtained. The results show that the approach is a sensible option and provide several hints for future improvement of the system.

Keywords: swarm intelligence, graph, stochastic sequencing, sequencing adaptation, ant colony optimization.

1 Introduction

Web based education has seen in recent years a significant increase in both its functionality as well as possible scenarios. Many research efforts are trying to take advantage of two fundamental characteristics of the Internet: small delays in communications (mostly independent of physical location) and big number of users. Social systems try to emulate the behaviour of social groups in real life. They try to extract some information of the behaviour of a group of students and use it to get some benefit for them (e.g. a better learning path, a better selection of materials, etc). In other words, they take advantage of the interactions between the different members of the learning community to help each of its members.
Social swarm systems are the result of applying swarm intelligence (henceforth SI) techniques to problems in which the elements of the systems studied are people. This paper presents an example of such a system. An adaptation of the ACO technique for its application to the sequencing of learning activities is presented and analyzed.
Paraschool is the French leading e-learning company, giving service to a vast number of students. They were 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 can apply and show interesting properties: adaptability and robustness. Two papers [2, 3] have already described the first steps toward the use of ACO-like algorithms with humans in the e-learning domain.

2 The man-hill paradigm

The Paraschool system divides the activities to be delivered to the student into courses and chapters. Courses can range from a short training course (e.g. a course on security when using heavy machinery) to a full academic year at a school (e.g. fourth grade class). Courses are divided into chapters. Inside each chapter, a graph of activities is defined. Every node in this graph is an activity, typically composed of a theory web page, an exercise that illustrates the concepts presented, and a final page that corrects the answer of the student and offers her some feedback. Edges of the graph represent possible transitions after a node has been covered. Nodes are not necessarily connected to themselves.
Probabilities are associated to every arc. These probabilities determine which will be the next learning unit to be delivered to the student. Units with a higher probability will be chosen more frequently than low probability ones. These probabilities are initially set by the pedagogical team of teachers that design the probability distribution of the graph in the form of pedagogical weights. These weights are normalized and then assigned to the arcs. Arcs that connect nodes are equivalent to restrictions in the sequencing of learning units. Probabilities associated to those arcs determine which transitions will usually be followed after each particular node. Taking into account both things, restrictions and probabilities, the sequencing of the students is specified stochastically, as can be seen in Figure 1.
In order to get some adaptation of the probabilities to the students, every student traversing the graph acts like an ant when interacting with the exercises and traversing the graph. After being succesful at an exercises, the student deposits positive pheromones (ϕ^+) behind; in the case of

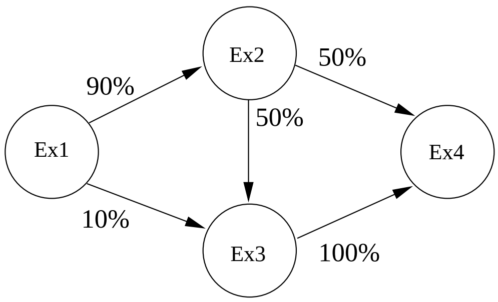

Fig. 1. Stochastic sequencing

success, negative pheromones (ϕ^-) are deposited. The former increase the probability of an arc being chosen, while the latter lower it.
Preliminar results with the system showed that the system was optimizing the sequences of exercises followed by the students towards "success": high marks and short paths. This meant that most difficult exercises were avoided, and so were comprehensive paths that might be defined by the pedagogical team. This behaviour was not positive for the students' learning, so the system had to be modified.
The Paraschool system was mainly based on learning through problem resolution. Taking into account that student learn more when confronted with challenges that have a difficulty level similar to their own knowledge [5], it was decided to modify the system to find paths where items would be moderately hard to validate at each interaction.
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 would not be discouraged by failing too often. An additional refinement was suggested: when the amount of pheromones on the arcs were not important enough to bear significant information, Paraschool asked that the relative pedagogical weight set by the teachers had more importance. Additionally, a personal pheromone (ϕ^P) prevented the same exercise to be repeated twice before a week had passed. The mathematical background is ommitted for the sake of space, but can be found in [4, 1].

3 Experimental scenario

We analyzed the data stored by the system during twenty-one months, from January 2005 to September 2006. The analysis of the empirical data presented serious difficulties. Due to logistic difficulties (many of the users of the systems were scattered around the country, interacting individually with the system) and differences of opinion with the company, there were not any pre-test information, only individual items regarding events that had happened at a specific point in time. For example, the Paraschool system might record that a student has solved exercise 245 at a specific date and time. There was not any common reference to compare the students at two specific points in time, in order to observe the evolution of the group.
In order to evaluate the effect of the system on the evolution of the students, it was necessary to divide the students population in at least two groups. The separation was operated with respect to the types of navigation. There are three main modes of navigation in the system: free, guided by a teacher and suggested by the man-hill.

Free navigation (F). The students have a table of contents with links to all the exercises in the courses in which they are registered. They can select any exercise and try it. When they do, a new arc might be created if it did not exist between the last exercise they tried and the current one. This type of transition amounted to approximately 50% of the total.

Table 1. Average of the evolution factor for each group (numbers represent points over 100)

Type	No filter	1 day	1 week
A	18.21	7.65	6.29
F	26.16	9.14	7.44
G	35.85	8.46	5.65

Guided navigation with a teacher (G). This mode is used when the Paraschool system is used in a blended learning environment as a complement to a traditional class. The teacher has all the information about the exercises, and instructs the students to try them at specific moments that fits in the normal course planning. These transitions represented 25% of the total.

Following of suggestions by the ant system (A). Every time the student finishes an exercise, the system analyses the outgoing arcs, calculates a fitness value for each one and selects several arcs using a stochastic contest [3]. These options are presented to the student as suggestions for the next exercise. When the student chooses one of these options, that is registered as an A transition. Approximately 25% of the total transitions were of type A.

It should be noted that this is not a taxonomy of students. Students are allowed to use any type of transitions they want. For a example, a student may follow the suggestions of the ant system for a sequence of three exercises, then go to the table of contents and select another exercise (free navigation), then follow the suggestions of the ants again. There is no limitation to this.

Results and discussion

The evolution of the students in each group was averaged. Different time filters are used, in order to avoid that too many repetitions of the same exercise are examined in a short time. The results are summarized in Table 1. Each row represents one group: group A mainly follows the suggestions of the ant system, group F navigates freely and group G executes the exercises according to the guidance of a teacher. Each column represents the time filter used: on the first column no filter is used, on the second column only those repetitions of an exercise that were separated at least 24h were considered, etc.

As the minimum time between repetitions increases, a logical decrease on the results due to forgetfulness for the three groups can be observed: the three obtained the highest increases when no time filter was applied, while the worst results are obtained with the filter of one week. However, the effect of time is different for the three groups. The results of group A diminish the less with the one day filter, while those of group G show the steepest decrease in both cases. This is summarized Table 2.

It can be observed that the group that followed the suggestions of the ant system (A) had much worse results than the other two groups when no

Table 2. Decrease on results due to forgetfulness

Type	No filter → 1 day	1 day → 1 week
A	57.99%	17.78%
F	65.06%	18.60%
G	76.40%	33.22%

time filter was used, but this difference diminished as longer time filters were applied. When the time filter is one week long, the differences are small with the other two groups (group G performs worse if the one week filter is used).

The analysis shows that preventing students of repeating the same exercise twice before one week had passed had a bad impact on their results. This restriction was present in the ant system by means of the personal pheromone ϕ^p. Students of groups F and G could —and did— repeat exercises in shorter periods of time, and they obtained better results.

Students of group A could repeat exercises before one week had passed, but they usually did not because they followed the suggestions of the system. The ant system rarely suggests a student to repeat the same exercise before one week has passed, as the multiplicative factor ϕ^p lowers significantly the fitness of that arc. The possibilities for a student from group A to repeat an exercise in a shorter time are either free navigation, or finding another arc (with a different ϕ^p) that led to the same exercise from a different place. Results suggest than a smaller repetition time would have allowed group A to perform better.

The difference was specially evident when no time filter was used: this showed that the hypothesis that students repeat exercises in a short period of time with high increases in scores but modest effects on their learning was true. Therefore, the existence of a restriction that prevented or restricted very early repetitions was a sensible option. However, the current configuration is excessive and a shorter time barrier would be more adequate.

On the other hand, it was expected that this "early repeat" behaviour would be specially evident on group F, but it was group G that showed a bigger dependency with the minimum interval between repetitions. Group G obtained also the best results when no time filter was applied. Both evidences suggest that students under guidance of a teacher repeated the same exercise more frequently than the other groups.

Without forgetting that ϕ^p should be better calibrated and allow for earlier repetitions of exercises, the analysis shows that the results of group A are comparable to the other two groups when a time filter of one week is used. Group F performed slightly better, and group G slightly worse, but the difference is not significant. This means that this ant-based approach for finding learning sequencings of learning activities in a auto-organizative fashion does not impose a negative burden on learning. This was specially important for the company. It is true that the results are not better than in the case of free navigation, but the overall result after this research is a better system. The system provides a set of exercises, with both self-directed navigation and sequencing suggestion modes. Every student has the freedom of selecting the sequencing strategy that he/she prefers and the effect on his/her learning will be similar.

4 Conclusions and future work

An approach for the automatic selection of good learning paths for students has been presented. It is based on an adaptation of the Ant Colony Optimization paradigm that has been used succesfully for different applications. The crucial differences between the behaviour of ants and that of people make it necessary to make changes to the paradigm.
The new approach was applied to an e-learning system that provided students with a sequence of exercises. The number of students and exercises made it feasible to adapt the ACO paradigm as follows. When the students interact with the system, they deposit pheromones behind them. Positive pheromones are deposited if exercises are solved correctly, reinforcing the path followed; when they fail, negative pheromones are deposited for the opposed effect. An additional level of adaptation was pursued using personal pheromones.
The results show that students that followed the suggestions of our system obtained similar results to those that preferred to do exercises freely, accessing them through a table of contents. However, the 1-week period showed itself to be excessive. This will be corrected in later versions.
Another interesting result is that students that solved the exercises following the indication of a teacher (in a blended learning scenario) obtained worse results in several cases. The reason for this is still unknown, and demands further investigation.

Acknowledgements This work was supported by project Mosaic Learning TSI2005-08225-C07 and grant BES-2003-0898.

References

1. S. Gutiérrez, G. Valigiani, Y. Yamont, P. Collet, and C. D. Kloos. A swarm intelligence approach for automatic auditing of pedagogical planning. In *The 7th IEEE International Conference on Advanced Learning Technologies (to appear)*. Springer-Verlag, July 2007.
2. Y. Semet, E. Lutton, and P. Collet. Ant colony optimisation for e-learning: Observing the emergence of pedagogical suggestions. In *IEEE Swarm Intelligence Symposium*, 2003.
3. Y. Semet, Y. Yamont, R. Biojout, E. Luton, and P. Collet. Artificial ant colonies and e-learning: An optimisation of pedagogical paths. In *10th International Conference on Human-Computer Interaction*, 2003.
4. G. Valigiani. *Développement d'un paradigme d'Optimisation par Hommilière et application a l'Enseignement Assisté par Ordinateur sur Internet (PhD thesis)*. Université du Litoral Côte d'Opale, 2006.
5. M. Wolfe, M. Schreiner, B. Rehder, D. Laham, P. Foltz, W. Kintsch, and T. Landauer. Learning from text: Matching readers and texts by latent semantic analysis. *Discourse Processes*, 25, 1998.