Multi-Agent Physical A* Using Large Pheromones

Ariel Felner, Yaron Shoshani, Israel A. Wagner and Alfred M. Bruckstein
Department of Computer Science
Bar Ilan University
Ramat Gan, Israel
{felner,shoshai}@cs.biu.ac.il

1. Introduction

This paper introduces the notion of *large pheromones* as a model of communication and global knowledge sharing in multi-agent systems. With *pheromones*, agents communicate by writing and reading data at the nodes of the graph that constitutes their environment [2]. Unlike ordinary pheromones where only a limited amount of data can be written in each node, in the *large pheromones* model, there is no restriction on the amount of data that can be written in the the nodes and thus each agent can write its entire knowledge in a node. We apply this model of communication to the multi-agent physical A* algorithm (MAPHA*) which is the multi agent version of Physical-A* (PHA*) [1]. These algorithms modify the A* algorithm to find shortest paths in physical environments with mobile agents that move around the environment and explore unknown territories. These algorithms are designed to minimize the travel effort of the agents.

We will show that increasing the amount of data that can be stored at each of the pheromones dramatically reduces the travel effort of the agents. With maximal usage of this model with unlimited size of pheromones the behavior of a multi-agent system is almost as good as with complete knowledge sharing between the agents.

2. Physical A*

Physical A* (PHA*) modifies the well-known A* to find the shortest path in graphs which correspond to a real physical environment. Consider a mobile agent who needs to find a shortest path between two physical locations and assume that only a very small portion of the environment graph is known to the agent. The mobile agent needs to activate the A* algorithm on this physical graph. Unlike traditional A*, for this type of graph, we cannot assume that expanding a node from the open list takes constant time. Here, to expand a node that is not known in advance, a mobile agent must first travel to that node in order to explore it and learn about its neighbors. The cost of the search in this case is the cost of moving an agent in a physical environment, i.e., it is proportional to the distance traveled by the agent. PHA* expands all the mandatory nodes that A* would expand and returns the shortest path between the two points but is designed to minimize the traveling effort of the agent by intelligently choosing the next assignment of the traveling agent.

PHA* works in two levels. The high level acts like a regular A* search algorithm: at each cycle it chooses the best node from the open-list for expansion. Nodes are evaluated according to a cost function $f(n) = g(n) + (h(n)$, If the node chosen by the high level has not been explored by the agent, the low level, which is a navigation algorithm, is activated to bring the agent to that node and explore it.

PHA* was generalized to the Multi-agent Physical A* (MAPHA*) where a number of agents cooperate in order to find the shortest path. The task is that these agents should explore the necessary portion of the graph, i.e., the A* nodes as fast as possible. The assumption in [1] was that each agent can communicate freely with all the other agents and share data at any time. Thus any information gathered by one agent is available and known to all of the other agents. MAPHA* also uses a two level framework. The high level chooses which nodes to expand, while the low level distributes the agents to navigate to these nodes.

2.1. Pheromones

A famous and interesting model for communicating is that of ant-robotics [2]. In this model, information is spread to other agents via *pheromones*, i.e., small amounts of data that are written by an agent at various places in the environment (e.g. nodes in the graph), and can be later used or modified by other agents visiting that node. Usually, it is assumed that a very small amount of data (no more than a few bytes) can be written in each pheromone.

We suggest a new model of communication which we call *large pheromones*. Unlike conventional (small) pheromones, we assume that an agent can write its entire knowledge base (e.g. a complete list of nodes and edges known to that agent) at each of the nodes. With today's hardware capabilities and computer architecture this is a reasonable assumption. For example, it is not un realistic to assume one megabyte of memory at a node which can store a graph of tens of thousands of nodes. We can also assume that the time to read and write data from the large pheromones can be omitted when considering the traveling time of the agents. This additional data storage in the *large pheromones* paradigm can help the agent to solve the problem faster and much more efficiently.

2.2. Spreading the data

With such large capacities of memory in each node, we present the following communication paradigm between the agents. Each agent has a partial knowledge of the entire environment. It maintains a database with a partial graph that is known to it. Similarly, each node holds a database with a partial graph that is 'known' to it, i.e., knowledge that was written to it by the agents. Whenever an agent reaches a node, it merges the data known to it with the data that is written in that node. The agent then writes the combined data in that node and updates its own database according to the new data that was obtained. We call this the *data-merge* operation. In this way data will be spread out very fast as long as agents visit many nodes in many areas of the graph and perform *data-merge* operations at all the nodes that they visit. We have modified the MAPHA* algorithm to use the large pheromones paradigm. At each node that a navigating agent visits, it performs a *data-merge* operation. This might have a positive effect on the A* that is activated by this agent as it knows more knowledge about the graph.

3. Experiments

We have implemented MAPHA* with large pheromones and performed many experiments on with different variations of this algorithm.

Figure 1 illustrates the time elapsed as a function of the number of agents that were used to find the shortest path with different versions of MAPHA* on Delaunay graph with 500 nodes. Since we assume constant speed, the time is reported as the distance traveled by the agents until the solution was found. There are 6 curves in the figure. The bottom curve corresponds to the best version of MAPHA* with full communication, and is used as a benchmark. Other curves show the overall time cost of different versions of MAPHA* with large pheromones. The top two curves show the behavior of $MAPHA*_{LP}$ where we only assumed a limited memory capacity at the nodes. In particular, in the top curve. 10 nodes were allowed to be written and the second curve allowed 20 nodes to be written at each pheromone. The rest

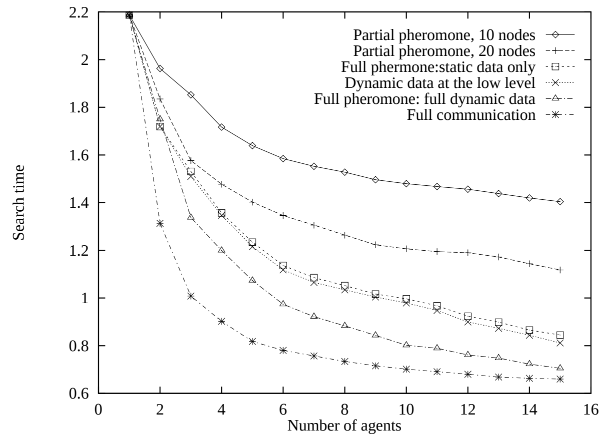

Figure 1. Foraging Productivity Results

of the curves assume unlimited data capacity and thus the entire graph can be written at each node. They differ in the amount of knowledge that was allowed to be written.

The results show a clear phenomenon. As the pheromone becomes larger and includes more knowledge a significant improvement in the overall time is obtained. This parametric effect is achieved even though no explicit control is forced by a centralized supervisor. The figure clearly shows that all the versions that use the large pheromones paradigm with unlimited memory capacity keep most of the potential of the full knowledge sharing paradigm. Their performance is rather close to the performance of the full communication version. This means that with large pheromones, data is spread to other agents rather fast and it is almost as good as full communication and full knowledge sharing. We have experimented with other sizes of graphs and on sparse and dense graphs and obtained similar results.

4. Conclusions

We believe that large pheromones are needed in any domain where global data from different areas of the environment is critical to the decision making of all agents at all times. In that case smaller pheromones will not do the job. Our problem of activating A* in a physical environment is an example for this. Since A* expands nodes in a global best-first search order local data is not enough for this as shown in figure 1.

References

[1] A. Felner, R. Stern, A. Ben-Yair, S. Kraus, and N. Netanyahu. Pha*: Finding the shortest path with a* in unknown physical environments. *Submited to JAIR*, 2003.

[2] A. Wagner and A. M. Bruckstein. ANTS: Agents, networks, trees, and subgraphs. *Future Generation Computer Systems Journal*, 16(8):915–926, 2000.