

Non-Optimal Multi-Agent Pathfinding Is Solved (Since 1984)

Gabriele Röger and Malte Helmert

University of Basel, Switzerland
{gabriele.roeger,malte.helmert}@unibas.ch

Abstract

Optimal solutions for multi-agent pathfinding problems are often too expensive to compute. For this reason, suboptimal approaches have been widely studied in the literature. Specifically, in recent years a number of efficient suboptimal algorithms that are complete for certain subclasses have been proposed at highly-rated robotics and AI conferences. However, it turns out that the problem of non-optimal multi-agent pathfinding has already been completely solved in another research community in the 1980s. In this paper, we would like to bring this earlier related work to the attention of the robotics and AI communities.

Introduction

Multi-agent pathfinding (MAPF) is an interesting problem which is relevant for many real-life applications like warehouse management, logistics or computer games.

A MAPF problem is given by a graph where each node either contains an agent or is unoccupied. An agent can move to an adjacent node if this node is unoccupied. The goal specification defines a goal node for each agent and the aim is to move all agents to their respective goal. There are two versions for this problem that have different optimization criteria: In the “sequential” variant we want to minimize the number of agent movements, whereas in the “parallel” version we would like to minimize the number of steps, where agents can move in parallel in each step if their destination nodes are unoccupied. Since in this work we are only interested in non-optimal approaches and we are not concerned with quality guarantees relative to the optimal solution cost, we do not need to distinguish these versions but refer only to MAPF problems in general.

Work in AI and Robotics

Most suboptimal MAPF techniques are incomplete decentralized methods. However, in recent years there have been several proposals of efficient suboptimal algorithms that are complete for certain subclasses of the problem.

Peasgood, Clark, and McPhee (2008) present an algorithm which is complete for trees with fewer agents than leaves in the tree. Khorshid, Holte, and Sturtevant (2011) propose a more general approach guaranteeing completeness for a larger class of trees.

The MAPP algorithm by Wang and Botea (2011) covers a fragment of general graphs, called *SLIDABLE*, whose characterization is defined in terms of the paths from the initial to the goal node of all agents. However, Khorshid, Holte, and Sturtevant show how each *SLIDABLE* instance induces a tree on which their algorithm can solve the original instance. Extended versions of MAPP can solve larger subclasses (Wang and Botea 2011), but their characterization is too involved for a brief presentation within the scope of this paper.

Another very recent approach, called *Push and Swap*, by Luna and Bekris (2011a; 2011b) is complete for a very large subclass of MAPF which is also easy to define, namely all instances with at least two unoccupied nodes.

The authors of these papers identify non-optimal MAPF as an interesting open problem. For example, Khorshid, Holte, and Sturtevant write that their “work is just one step in classifying problems which can be solved in polynomial time”. Wang and Botea emphasize as a contribution that their approach “identifies classes of multi-agent path planning problems that can be solved in polynomial time”. Finally, Luna and Bekris state that “in comparison to existing alternatives that provide completeness guarantees for certain problem subclasses, the proposed method provides similar guarantees for a much wider problem class” (2011b).

The Solution

Given that the work cited above has been published at major AI and robotics conferences and journals, we have the impression that not only the authors but also these research communities in general are not aware that non-optimal multi-agent pathfinding has been fully solved in 1984 under the name of *pebble motion on graphs*. Hence, with this paper, we would like to bring the work by Wilson (1974) and by Kornhauser, Miller, and Spirakis (1984) to attention.

This work builds on the observation that for a graph with n nodes one can easily, with less than n^2 moves, move all agents to nodes that must also be occupied by some agent in the goal configuration. The resulting configuration together with the goal configuration induces a permutation on the occupied nodes which corresponds to a solution to the problem. Based on well-known results on permutation groups, Wilson shows that for biconnected graphs G with only one unoccupied node (like the 15-puzzle), the problem is solvable iff the single biconnected component of G is not bipar-




Figure 1: Exception θ_0 to Wilson's criterion.

tite or the induced permutation is even. The only exceptions are polygons and the graph θ_0 shown in Figure 1. For polygons, one needs to check whether the order of the agents is preserved. Graph θ_0 is a small special case that can be solved with a lookup table.

From the proof given by Wilson it is possible to derive a solution for every solvable MAPF problem with $n - 1$ agents with $O(n^5)$ moves (Kornhauser 1984).

Kornhauser, Miller, and Spirakis (1984) modify Wilson's proof for problems with one unoccupied node to derive solutions with fewer moves and present a polynomial-time algorithm that is complete for the full class of MAPF problems, allowing instances with fewer than $n - 1$ agents as well as separable (non-biconnected) graphs.

The algorithm decomposes the problem into subproblems, each defined by a set of agents which can reach the same set of nodes and the subgraph induced by these nodes. The original problem is solvable iff each of these subproblems is solvable. For solvability, the goal node for each agent must lie in the same subproblem as the initial node. If this test passes, problems with $n - 1$ agents are solvable if each subproblem is solvable according to Wilson's criterion. With fewer agents, every subproblem that is not a polygon is solvable. Polygons are solvable if the start and end nodes of the agents differ only by a cyclic permutation.

Kornhauser, Miller, and Spirakis furthermore show how to efficiently generate a solution with $O(n^3)$ moves for solvable problems. They also show that this is in a certain sense tight, giving a class of problems requiring $O(n^3)$ moves.

Discussion

One reason why the work by Kornhauser, Miller, and Spirakis fell into oblivion could be that the only archival publication is very sketchy and often refers to a "final version" which never appeared. However, all results are described in detail in Kornhauser's master's thesis, which is available as technical report (Kornhauser 1984).

The details reveal that many recent findings and concepts have already been covered in these old papers. For example, Khorshid, Holte, and Sturtevant specify three conditions which are sufficient for a tree to be solvable for any configuration of agents. These are a special case of Kornhauser's criteria when considering only trees, but Kornhauser gives a more precise analysis, specifying sufficient and *necessary* conditions (also considering the actual configuration).

The approach by Peasgood, Clark, and McPhee and the Push and Swap algorithm by Luna and Bekris also rely on observations that were already contained in Kornhauser's thesis. These algorithms can possibly be interpreted as instantiations of the general algorithm by Kornhauser for the

case with at least two unoccupied nodes.

This somewhat cautious statement ("possibly") highlights another reason why the work by Kornhauser, Miller, and Spirakis is not more widely known: the approach is not described algorithmically, but must be derived from a number of proofs in the paper. This requires significant effort, and to the best of our knowledge neither Wilson's result nor the one by Kornhauser, Miller, and Spirakis have ever been implemented. However, the description of the approach in terms of group-theoretic proofs also has advantages: it provides a deeper understanding of the problem and is rather generic, leaving a lot of choices open. The stated guarantees hold independently of the actual instantiation, but we expect that the decisions still have a significant impact on the solution quality, leaving room for interesting optimizations.

For this reason, we are currently deriving an algorithmic description of the approach and are working on an implementation. The aim of this future work is on the one hand a proper comparison to the recent algorithms – experimentally as well as in terms of a clear theoretical demarcation. On the other hand, we would like to make the approach more easily accessible to the community.

References

- Khorshid, M.; Holte, R. C.; and Sturtevant, N. 2011. A polynomial-time algorithm for non-optimal multi-agent pathfinding. In Borrajo, D.; Likhachev, M.; and López, C. L., eds., *Proceedings of the Fourth Annual Symposium on Combinatorial Search (SOCS 2011)*, 76–83. AAAI Press.
- Kornhauser, D.; Miller, G. L.; and Spirakis, P. 1984. Coordinating pebble motion on graphs, the diameter of permutation groups, and applications. In *Proceedings of the 25th Annual Symposium on Foundations of Computer Science (FOCS 1984)*, 241–250.
- Kornhauser, D. M. 1984. Coordinating pebble motion on graphs, the diameter of permutation groups, and applications. Technical Report MIT/LCS/TR-320, Massachusetts Institute of Technology.
- Luna, R., and Bekris, K. E. 2011a. Efficient and complete centralized multi-robot path planning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2011)*, 3268–3275.
- Luna, R., and Bekris, K. E. 2011b. Push and swap: Fast cooperative path-finding with completeness guarantees. In Walsh, T., ed., *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI'11)*, 294–300.
- Peasgood, M.; Clark, C. M.; and McPhee, J. 2008. Complete and scalable multi-robot roadmap coordination planning. *IEEE Transactions on Robotics* 24(2):283–292.
- Wang, K.-H. C., and Botea, A. 2011. MAPP: a scalable multi-agent path planning algorithm with tractability and completeness guarantees. *Journal of Artificial Intelligence Research* 42:55–90.
- Wilson, R. M. 1974. Graph puzzles, homotopy, and the alternating group. *Journal of Combinatorial Theory (Series B)* 16(1):86–96.