A COMPARISON OF DIPLOMACY GAMEBOARD GRAPH SEARCH ALGORITHMS

Daniel P. Stormont and Vicki H. Allan
Department of Computer Science, Utah State University, Logan, Utah, U.S.A.

Keywords: Agent, Diplomacy, Search.

Abstract: The boardgame *Diplomacy* has been used as a testbed for multiagent systems almost since the time of its introduction in 1959. The reason is that the game presents a number of interesting challenges to artificial intelligence researchers: a state space that is too large to be tackled by brute force searches, imperfect information due to simultaneous movement, no random elements, and non-binding negotiations between the seven players. This paper looks at just one aspect of creating an agent for playing *Diplomacy* – finding the fewest number of moves to achieve a victory in the game, if the player was unopposed. This planning function forms the basis for a more sophisticated move planner that also takes into account the game state and the other players. Three search algorithms are compared to determine which is the most effective (in terms of the number of map nodes expanded during the search).

1 INTRODUCTION

Researchers in multiagent systems seek environments that provide difficult algorithmic challenges and realistic situations in order to advance the state of the art. The boardgame *Diplomacy* is just such an environment. It provides challenges in many areas of active research for agents: planning, cooperation, negotiation, trust, and coalition formation, just to name a few. It is for this reason that there are a number of testbeds for multiagent systems based on *Diplomacy*, including the *Diplomacy* Artificial Intelligence Development Environment (DAIDE) (DAIDE, 2011) and DipGame (Fabregues and Sierra, 2009).

This paper addresses one element of creating a planner for a *Diplomacy* agent: an efficient search algorithm for determining the shortest path to achieving victory in the game. For the purposes of this paper, the search is run from the starting positions for each of the seven players and the optimal path to victory is determined without taking into account the position of any opposing players or other elements that need to be considered for an optimal planner; such as negotiated agreements between players, the relative strengths of the players, and the need for cooperation between agents or units of a single agent in order to achieve the goals. Thus, the planner described here is just the basis of a more sophisticated planning agent and the search described in this paper would need to be rerun or updated as the game state changes during play.

After a brief introduction to the game *Diplomacy* with an emphasis on the gameboard and its representation, the paper will describe the three search algorithms selected for comparison, describe the design of the software developed for the comparison, detail the setup of the experiment, provide expected and actual results from the comparison, and provide conclusions and plans for future work on a *Diplomacy* agent based on this work.

2 THE GAME OF DIPLOMACY

The game of *Diplomacy* was created by Allan B. Calhamer in 1959, based on the work of one of his college professors, Sidney B. Fay (Calhamer, 2000) (Fay, 1934). The game attempts to recreate the political situation in Europe prior to the First World War and the system of secret alliances that initially maintained the peace in Europe, but would eventually lead to the greatest conflict in history up to that time. To do this, the game incorporates a number of elements into the game play. There is a defined negotiation period before every movement phase during a turn of the game. There are no rules for how these negotiations occur: they can be secret or public, at the gameboard or in

P. Stormont D. and H. Allan V..
A COMPARISON OF DIPLOMACY GAMEBOARD GRAPH SEARCH ALGORITHMS.
DOI: 10.5220/0003753103710374
In *Proceedings of the 4th International Conference on Agents and Artificial Intelligence* (ICAART-2012), pages 371-374
ISBN: 978-989-8425-96-6

another room, and any agreements reached are non-binding. The only restriction is that the negotiation period is typically time limited. After negotiations, all players submit their planned moves in writing (or via a computer) and the moves are adjudicated simultaneously. There are no elements of chance in the movements (no dice rolling or card drawing) and the initial strengths of the players will dictate cooperation between players, at least initially, as no player is strong enough to win the game without assistance.

2.1 Representation of the Gameboard

Figure 1 shows the standard *Diplomacy* gameboard. There are 75 provinces on the board (81 if coastal areas are taken into account), which are connected along their borders, complicating the connections between them. There are also three different types of provinces: coast, land, and sea. Land areas can only be occupied by armies, sea areas can only be occupied by fleets, and coastal areas can be occupied by either. The small circles on the map are supply centers. The objective of the game (and the goal of the search algorithms described in this paper) is for one player to control 16 supply centers by either having been the last player to pass through that province (it is actually a little more complex in the game since control of supply centers is determined on every other turn - that has been ignored for the purposes of this study) or to have a unit in the province.

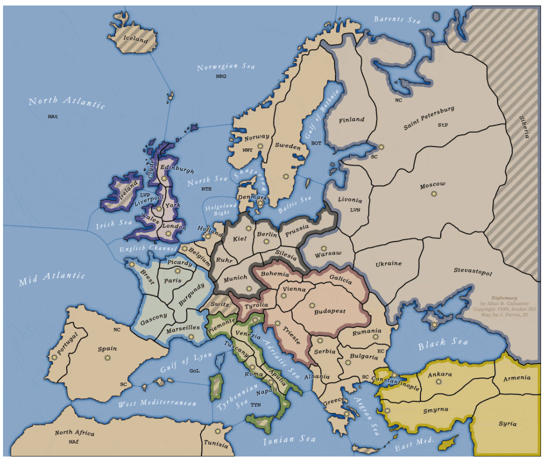

Figure 1: The *Diplomacy* gameboard.

Figure 2 shows the graph generated from the map in figure 1. In this graph, the provinces are shown as different colored nodes. Coast provinces are tan, land provinces are green, and sea provinces are blue. A unit can traverse any of the edges connecting nodes, as allowed by the unit type. For example, an army can not follow an edge from a coast region to a sea region since only fleets can occupy a sea region. This unit type-based limitation on movement is the reason for the three provinces with dashed lines (Bulgaria, Spain, and Saint Petersburg). These provinces have more than one coast region, which limits the possible traversals by a fleet from those provinces to a neighboring province. Provinces with supply centers are indicated by bold text in the graph.

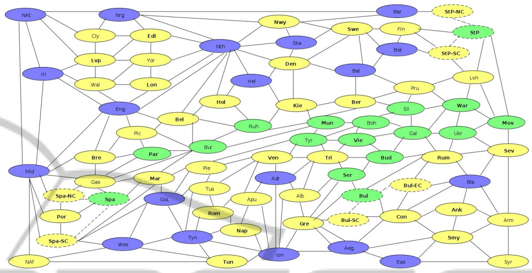

Figure 2: A graph showing the connections between provinces on the *Diplomacy* gameboard.

2.2 Starting Positions

Figure 3 shows the starting positions and unit types for each of the seven players in *Diplomacy*: Austria, England, France, Germany, Italy, Russia, and Turkey. Each of the players starts with three units, except for Russia, which starts with four. These starting positions and unit types are critical to this comparison of search algorithms as they comprise the first ply of the search tree for each of the players.

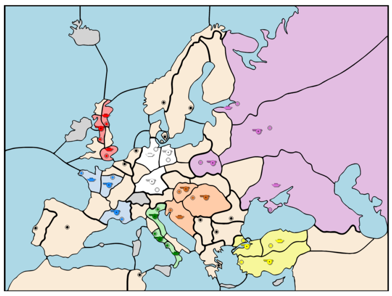

Figure 3: A *Diplomacy* gameboard showing the starting locations and unit types for each player.

3 SEARCH ALGORITHMS SELECTED

Three common graph search algorithms were selected for comparison in finding the shortest number of moves for each player to reach the goal state of having passed through or occupying 16 supply centers on the *Diplomacy* map graph shown in figure 2. For each player, the nodes occupied in the starting position are selected as the first ply of the graph search, as shown in figure 4. These starting positions also satisfy the goal condition for occupying three supply centers (four for Russia), meaning the algorithms have to find the shortest path(s) through the graph, starting with the first ply of the search in the player's home nation to pass through or occupy an additional 13 supply centers (12 for Russia).

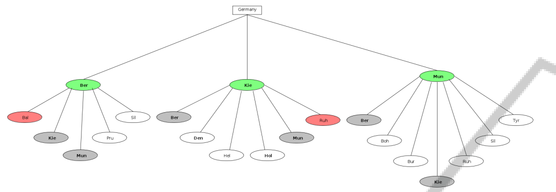

Figure 4: An example of the first ply of the *Diplomacy* gameboard search for the German player.

In figure 4, the root of the search tree is the nation for that player (Germany in the example). The first level of the search tree are the nodes corresponding to the starting positions for that player. (In the case of Germany, there are armies in Berlin and Munich and a fleet in Kiel.) Note that the nodes for the starting positions are colored green. That is because these nodes are supply centers that are occupied by the player at the beginning of the game. The next layer shows the expansion of the graph nodes for provinces connected to the starting provinces in the *Diplomacy* graph. Red nodes are provinces that can't be occupied by that type of unit. For example, the unit in Berlin can not enter the Baltic Sea node because it is an army and only fleets can occupy sea provinces. Gray nodes are nodes that are in the visited list because they have already been traversed. (The gray nodes in the figure correspond to the starting nodes.) Finally, nodes with bold text are supply centers (goal states). The criteria for selecting the search algorithms were that the search algorithm be usable on a highly connected graph, like the *Diplomacy* map graph, and that they use varying degrees of heuristics. One search was selected that doesn't utilize heuristic values at all (breadth first search), one was selected that only utilizes heuristic values (greedy best first search), and one was selected that utilizes both heuristic values and the actual cost of the path found so far (A* search). The details of these algorithms can be found in (Russell and Norvig, 2003). Brief descriptions of the algorithms and the heuristics applied are provided in the sections that follow.

3.1 Breadth First Search

Breadth first search was chosen as the non-heuristic search algorithm. In the example shown in figure 4, the algorithm will start by inserting the three starting provinces (Berlin, Kiel, and Munich) into a first-in/first-out (FIFO) queue. The goal count is also set to three, since the three starting positions are supply centers (goal nodes). Since they are goal nodes, when they are added to the queue, these provinces are added to the visited list to prevent cycles. Then, the provinces will be extracted from the queue in alphabetical order (since that is the order they were inserted in) and the provinces connected to them that are legal moves for the type of unit occupying the province and are not already in the visited list will be added to end of the queue. To take the province of Berlin as an example, the Baltic Sea will not be added to the queue because it would not be a legal move for the army in Berlin and Kiel and Munich will not be added to the queue because they are already in the visited list. This means Prussia and Silesia will be added to the queue and then Kiel will be removed from the queue for expansion. This continues until the supply center count reaches 16, at which time the search is terminated.

3.2 Greedy Best First Search

Greedy best first search was selected as an example of a heuristic-only search algorithm. Greedy best first search proceeds as described in the breadth first search, except that each node has an associated f-value where $f(x) = h(x)$. The heuristic value ($h(x)$) is set to five initially, for all players except Russia. Because Russia starts with four units instead of three, the initial heuristic value is set to four. These heuristic values reflect a theoretical minimum value for the distance to the number of nodes that unit needs to contribute to the overall goal of 16 supply centers. When a goal node is entered into the queue, the heuristic value for it is decremented by one, so in the case of Germany as illustrated in figure 4, the initial heuristic will be $h(x) = 5$. When Berlin, Kiel, and Munich are added to the queue, the heuristic value associated with each of these provinces will be decremented to $h(x) = 4$. The queue is a priority queue, so lower f-values will be extracted from the queue first in FIFO order. In this case, there will be no difference in the extraction order from the queue; however, when Denmark and Holland are added to the queue, their heuristic values will be $h(x) = 3$ because they are goal nodes, so they will be extracted from the queue before the remaining provinces on the same level of the tree – and their children will also have $h(x) = 3$, unless the child is also a supply center, in which case it will have a value of $h(x) = 2$. The search essentially becomes a depth first search on these nodes at this point.

3.3 A* Search

A* search was selected as an example of a search that uses both heuristic and actual cost values. The f-value for A* search is $f(x) = h(x) + g(x)$, where the $h(x)$ is the same as described in the previous section on greedy best first search, but $g(x)$ is a cost function that is incremented by one for every level of expansion of the search tree. Using the example of figure 4, the initial f-value is $f(x) = h(x) + g(x) = 5 + 0 = 5$. After the first three nodes are inserted into the priority queue (the same as for greedy best first search), the f-values are $f(x) = 4 + 1 = 5$ so they will still be extracted in FIFO order. When Denmark and Holland are inserted into the queue, their f-values will be $f(x) = 3 + 2 = 5$, but the f-values of the other provinces that are not supply centers (for example, Helgoland Bight) will be $f(x) = 4 + 2 = 6$, so these two provinces will be extracted from the queue first for expansion.

4 EXPERIMENTAL SETUP

To test the search algorithms, the three algorithms are used to search the graph of starting positions for each of the seven players in *Diplomacy*. The number of nodes expanded during the search for each of the nations and the total nodes expanded are used as the basis of comparison for the three algorithms. Because the starting positions never change, there is no reason to run the tests more than once for each algorithm.

5 RESULTS

The following sections summarize the results of the search algorithm tests for each of the three search algorithms. The theoretical results are a hand-calculated value of expanding the search tree for the German player, as was illustrated for the first two levels in figure 4. The actual results are the results of executing the code for each of the three search algorithms for all seven players.

5.1 Theoretical Results

Table 1 shows the result of manually exercising the three search algorithms on the starting positions shown in figure 4. It is somewhat surprising that, at least for the starting position of Germany, breadth first search expands fewer nodes in the search tree than either greedy best first search or A* search. This may be an aberration due to the unique starting position of Germany in what is essentially the center of the map, so other starting positions also need to be evaluated.

Table 1: Nodes Expanded by Search Algorithms for Starting Position of Germany.

BFS	Greedy BFS	A* Search
33	41	36

5.2 Actual Results

Implementation and testing of the search algorithms described here continues in simple agents that play a "no-press" (no negotiations) game of *Diplomacy*.

6 CONCLUSIONS AND FUTURE WORK

The search described in this study is just a preliminary planning step for a practical planning agent for playing the game of *Diplomacy*. A planning agent would also need to take into account the locations of other units on the gameboard (both hostile and friendly), the relative strengths of the other players, any negotiated agreements with other players, the need to protect territory already captured, the need to provide support to other units in order to capture territory, changes in the numbers of units on the board, retreats, and the changing state of the gameboard each turn.

REFERENCES

Calhamer, A. B. (2000). *Calhamer on Diplomacy: The Boardgame "Diplomacy" and Diplomatic History*. 1stBooksLibrary, Bloomington, Indiana.

DAIDE (2011). Reading.

Fabregues, A. and Sierra, C. (2009). A testbed for multiagent systems. Technical Report IIIA-TR-2009-09, Autonomous University of Barcelona.

Fay, S. B. (1934). *The Origins of the World War*. The Macmillan Company, New York.

Russell, S. and Norvig, P. (2003). *Artificial Intelligence: A Modern Approach*. Pearson Education, Upper Saddle River, New Jersey.