

Influence Graphs: a Technique for Evaluating the State of the World

Mikhail Utkin
University of Nevada, Reno

ABSTRACT

This paper gives a brief overview of influence maps - an existing method of game state evaluation in strategy board games, explains why this method cannot be applied to state evaluation of certain board games, and proposes a new method for state evaluation of these games – influence graphs. The new method is explained and its use is demonstrated in the game of Diplomacy. The paper finally cites examples of other games and real world applications where influence graphs can be used and looks at possible future improvements to this new method.

I. INTRODUCTION

Games are a significant driver of AI research. The abstract nature of games makes them an appealing subject for study [1]. Developing computer players for games (game AI) is a challenging and intriguing task. Such knowledge intensive approaches as finite state machines and rule-based systems are used to design intelligent game agents. These approaches require significant player and developer resources to create and tune to play competently and, thus, suffer from the knowledge acquisition bottleneck well known to AI researchers. In addition, these approaches work well until human players learn their habits and weaknesses [2]. The human players then find that the agents in games are unintelligent and predictable while they expect the agents to behave intelligently by being cunning, flexible, unpredictable, challenging to play against and able to adapt and vary their strategies and responses. Furthermore, players believe that agents' actions and reactions in games should demonstrate an awareness of events in their immediate surroundings. However, many games are proliferated with agents that do not demonstrate even a basic awareness of the situation around them [3].

For example, in the game *The Sims* the agents constantly receive information from the environment, and the AI is embedded in the objects in the environment, known as *Smart Terrain*. Each agent has various motivations and needs and each object in the terrain broadcasts how it can satisfy those needs. In *The Sims* the agents' behavior is autonomous and emergent, based on their current needs and their environment [3].

Another approach that is applicable to the problem of agents reacting to the game environment is a technique used in many strategy games, influence maps (IMs). As this paper will show, influence maps, while they are a useful tool, are not applicable to some board games. The problem is that not all game boards can be represented as a grid, which is a

requirement for the use of influence maps. This paper introduces a new method for game state evaluation – influence graphs – and demonstrates its use in Diplomacy, a strategy board game. Influence maps and their use are explained next.

II. INFLUENCE MAPS

Influence maps divide the game map into a grid of cells, with multiple layers of cells that each contains different information about the game world. The values for each cell in each layer are first calculated based on the current state of the game and then the values are propagated to nearby cells, spreading the influence of each cell. The value of an IM cell could be a summation of the natural resources present in that square, the distance to the closest enemy, or the number of friendly units in the vicinity [2]. Currently, influence maps are used in games for strategic, high-level decision making. However, it would also be possible to use them for tactical, low-level decision-making, such as individual agents or units reacting to the environment [3].

Several IMs are created and then combined to form the spatial decision making system. For example, two influence maps may be created, the first using an IM function which produces high values near vulnerable enemies, the second IM function producing high negative values near powerful enemies. Then those two influence maps are combined via a weighted sum. High valued points in the IM resulting from the summation are good places to attack - places where you can strike vulnerable enemies while avoiding powerful ones. The final step is to analyze the resultant IM and translate it into orders which can be assigned to units. In this example the highest valued point is taken, and the troops are told to attack there [2].

The advantage of influence maps over methods that are currently used in games, such as *Smart Terrain* in *The Sims*, is that the agent is presented with a single value (calculated using the weighted sum to combine all the factors) instead of numerous messages being sent to the agent about the environment [3].

Figure 1 shows an example of a result of combining several influence maps [2]. In this example, the two triangles represent two small boats equipped with rocket propelled grenade launchers. Their task is to attack an oil platform - pentagon, which is being guarded by a destroyer – hexagon. Each unit in this system has an influence map associated with it. Each small boat adds a circle of influence to its IM increasing the values assigned to the squares within the circle. It is up to the designer of the system to decide what the radius of the circle should be and by how much the values of the squares within a circle should be increased. Let's say that in

this case the radius of a circle shows the areas that can be reached by a shot from a grenade launcher installed on a small boat.

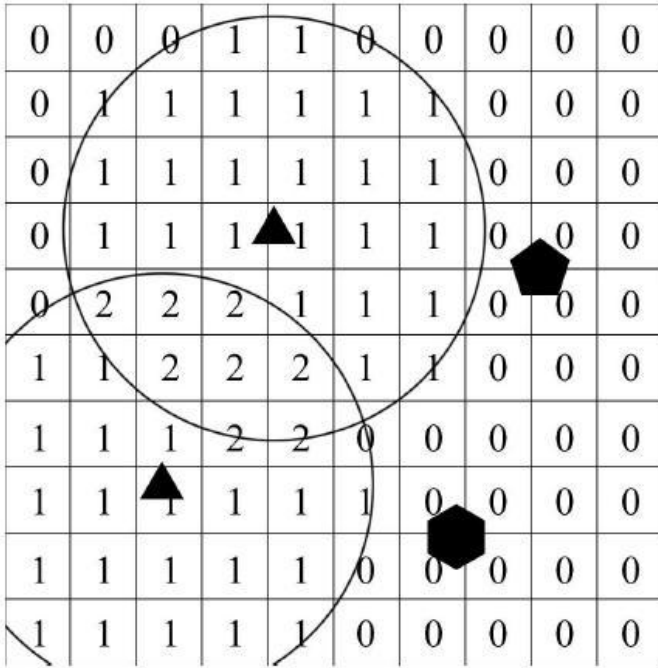


Figure 1: Combination of influence maps

When all influence maps are combined, the values of the square (cell) in individual IMs are added, and the squares in the area where the two circles intersect have a value of 2 assigned to them. The implication here is that these squares are especially vulnerable because they are within reach of grenade launchers from both small boats [2].

The game of Diplomacy, a strategy board game, shares a type of decisions with other strategy games. These decisions can be characterized as special reasoning problems: which parts of the world should a player control, which defensive installations to assault or how to outmaneuver an opponent in a battle. Let us see if and how influence maps can be used in Diplomacy. But first, a relatively detailed description of the game is presented.

III. THE GAME OF DIPLOMACY

Diplomacy was invented by Allan B. Calhamer, an American graduate student of history, political geography and law, all of which disciplines served him well in perfecting his game. Diplomacy received careful testing and constant revision before being marketed in its modern form. The idea which began to take shape in Allan Calhamer’s mind as early as 1945 did not in fact reach the public until 1959, by which time it had been polished and refined into a superbly balanced game [4]. Diplomacy is a seven-player game played on a board representing pre-World War I Europe (Figure 2).

Each player controls one of the seven Great Powers – Britain, France, Germany, Italy, Russia, Austria-Hungary, and Turkey, which are represented on the board by their armies and fleets. The object of the game is control of Europe.

The board is split into provinces, either inland, coastal or sea, which can only be occupied by a single unit at a time. An army can occupy either an inland or a coastal province. A fleet can occupy either a coastal province or a sea province. Thirty-four land provinces have one Supply Center (SC) each.

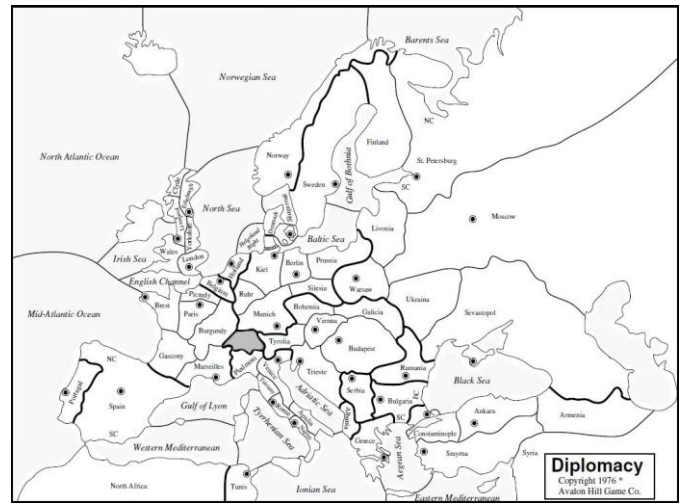


Figure 2: Diplomacy map, The Avalon Hill Game Co. 1976

Control of each center entitles the owner to a single unit. Each power begins with three supply centers (apart from Russia which has four) and the rest start as neutral. Each year is split into Spring and Fall turns, starting with Spring 1901. In each turn each player can negotiate with their rivals and then give orders to all his or her units. These orders are simultaneously revealed and processed. All units that are dislodged by an enemy attack must retreat and are given orders which are then revealed and processed. At the end of each Fall turn all occupied centers fall under the control of the player occupying them, and unoccupied centers remain under the control of the last player to have occupied them at the end of a previous Fall turn. Then each player counts the number of units and Supply Centers that he or she controls and builds or disbands units to ensure that he or she has no more units than Supply Centers. Units can only be built in a power’s original (‘home’) Supply Centers. If one player controls eighteen centers at the end of a Fall turn, then they win. The assumption is that the sixteen centers belonging to the other players cannot stop the eventual conquest of Europe. Crucially, only one unit can occupy a province at a time, and all units have equal strength. Units can move, hold or support a move of another unit.

Supporting effectively transfers the strength of one unit to another, but only if the supporting unit itself is not attacked by another unit. If the supporting unit is attacked (by a unit that it is not supporting an attack on) then the support has no effect, and is described as being “cut”, even if the attacking unit is dislodged. If two or more units contest a space, then the one with the most support occupies it, dislodging the occupant. In Figure 3 the support from Silesia is not cut, because the attack is coming from the province that the unit is supporting a move against, so the army in Warsaw is dislodged by the army in Prussia.

If both units have equal strength then they standoff and “bounce”, and neither moves but any occupants of the space are not dislodged. Figure 4 shows a German army in Prussia moving to Warsaw, supported by an army in Silesia. But the army attacking Silesia from Bohemia cuts support the army in Silesia is providing, so everyone bounces.



Figure 3: Attack with support



Figure 4: Support being “cut”

Support cannot be refused, even if from a unit of another power and no unit can cause another unit belonging to the same power to be dislodged. Units cannot swap places, unless one is convoyed, but three or more units can move in a circle [5]. Fleets can take actions that are different from land armies,

but these differences are not essential for the purpose of this paper: it will use a modified version of Diplomacy without fleets.

There are many variants of Diplomacy, using different maps and/or some additional rules. The “no press” Diplomacy prohibits communications between players (Diplomacy without diplomacy). Players then have to act based on the state of the map and their opponents’ previous orders.

The usefulness of influence maps in the game of Diplomacy will be analyzed next, and it will be shown why this tool is not helpful.

IV. INAPPLICABILITY OF INFLUENCE MAPS TO DIPLOMACY

An influence map is a grid placed over the game world. One of definitions of a grid is: a network of horizontal and perpendicular lines, *uniformly spaced*, for locating points on a map, chart, building plan, or aerial photograph by means of a system of coordinates [6]. The *uniformly spaced* part of the definition makes representing the Diplomacy board as a grid impossible: uniform spacing implies that, in a general case, a cell of a grid has exactly four neighbors, making the grid a specialized case of a graph, where each vertex representing a cell, except the cells on the edges and in the corners of the grid, is connected to four other vertices. A province in Diplomacy can have from as few as two neighbors (Portugal bordering on Mid-Atlantic Ocean and Spain) to as many as eight (Burgundy bordering on Paris, Picardy, Belgium, Ruhr, Munich, Switzerland, Marseille, and Gascony).

A regularized Diplomacy map is presented in Figure 5 [7]. It is topologically equivalent to Avalon Hill’s map in Figure 2 and it is on a uniform grid with vertical and horizontal lines. All regions on this map are at least two units in size.

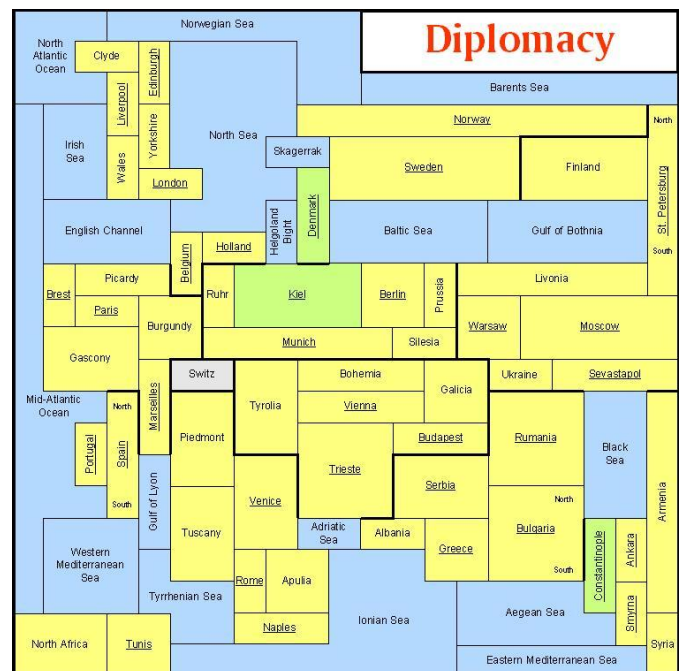


Figure 5: Regularized Diplomacy map

The Diplomacy map cannot be regularized so that each province occupies a single cell on a grid. At least some provinces will occupy at least two cells, as in the example in Figure 5. The fact that some regions occupy multiple grid cells makes influence maps inapplicable to the Diplomacy board. Influence maps have grid cells as their basic units. Even if influence maps are used on the normalized map of Diplomacy where regions are transformed to occupy several cells of a grid, as in the example in Figure 5, once the influence maps are combined, it is possible and, in fact, very likely that several cells covered by a single Diplomacy province will have different resulting values as shown in Figure 6.

4	3
3	2

Figure 6: The province of Berlin from Figure 5 covering four grid cells possibly having different values

This would contradict the basic feature of the game, which is that a Diplomacy province is a holistic and indivisible entity. The Diplomacy map in Figure 5 represents a more generalized version of a graph than a grid is. Thus, while the idea of influence may be used in Diplomacy, influence maps cannot be. A different method must be used for measuring influence on a Diplomacy map. This paper proposes a new method - influence graphs - for evaluating the state of the world which cannot be represented by a grid.

V. INFLUENCE GRAPHS

This paper provides a definition of an influence graph which is better understood when compared to the intensity of light coming from a light source and decreasing with distance. In this case, the light intensity is the ‘influence’ of the source: it is the greatest at the source and diminishes further away from the source. The ‘influence’ is quantified in numeric terms and is discrete in influence graphs. If the world is presented as a graph where objects are located at vertices, and the source of influence is an object located at a vertex, the influence of the source decreases equally along the paths of equal lengths from the source vertex to other vertices. The concept is illustrated by Figure 7. Here the source of influence located at the central vertex has influence of five which decreases by the same quantity (one) at all vertices the length of the path to which from the source vertex equals one.




Figure 7: Influence graph with diminishing influence

To calculate an object’s influence at vertices other than those adjacent to the vertex where the object is located, the shortest path from the object’s location to these other vertices needs to be identified. A number of alternative algorithms for identifying the shortest path are available. Probably the most efficient algorithm has been developed by Dijkstra [8]. As the game progresses and armies move on the board, any province could become a source of influence. The shortest path for each and every vertex representing a province should be calculated. It only needs to be done once at the beginning of the game. This information can then be stored and used to recalculate combined influence graphs after every set of moves. In a general case, if a Diplomacy map can be represented by a graph $G = (P, B)$ where P is a set of vertices (provinces) and B is the set of edges (borders between provinces) then the complexity of the algorithm for finding shortest paths for each vertex as a source is: $O(P^3)$. But if a binary heap is used in Dijkstra’s algorithm to store vertex labels then the complexity of the algorithm is $O(P \cdot (P+B) \cdot \log P)$ [9]. If the number of provinces is n and the number of objects which exert influence is m then the complexity of the algorithm for recalculating a combined influence graph after a set of moves is $O(n \cdot m)$. The task of identifying the shortest path is only made easier by the fact that a Diplomacy graph is not directed and all the edges have the same cost. The key point here is that the decrease in an object’s influence depends on the length of the path from the source vertex (the influencing object’s location) to the vertex where the influence is being calculated and that the decrease in influence is equal along any paths of the same length.

Influence graphs are a generalized case of influence maps. If influence graphs are applied to a game with a board that can be represented as a grid, the values for the influence of an object for all the cells in the grid can be calculated as described above for an influence graph.

		5		

Figure 8: A grid with a cell containing an object with an influence of five

3	3	3	3	3
3	4	4	4	3
3	4	5	4	3
3	4	4	4	3
3	3	3	3	3

Figure 9: A grid with calculated influence of the central cell at all other cells

For example, if a game board can be presented as a 5x5 grid with an object in the central cell having an influence of five

(Figure 8) then the object's influence in all other cells can be calculated using the shortest path algorithm if all cells are treated as vertices of a graph (Figure 9).

The term 'influence graph' has been previously used in the scientific community. It is appropriate at this point to look at other definitions of an influence graph and then compare them with the new definition which is applicable to board games.

Sphere of influence graphs (SIG) were first introduced by Toussaint as a type of proximity graph for use in pattern recognition, computer vision and other low-level vision tasks. A *random* sphere of influence graph (RSIG) is constructed as follows: Consider n points uniformly and independently distributed within the unit square in d dimensions. Around each point, X_i , draw an open ball ("sphere of influence") with radius equal to the distance to X_i 's nearest neighbor. Finally, draw an edge between two points if their spheres of influence intersect [10]. The concept is shown in Figure 10 [11].

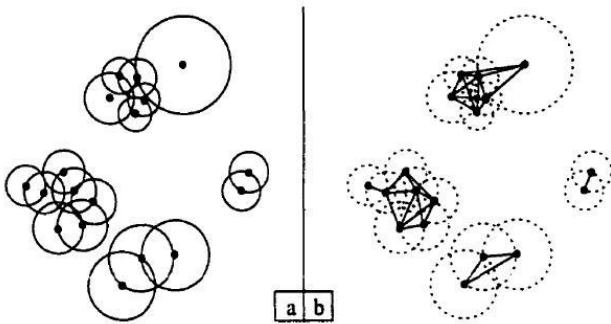


Figure 10: A set of points, (a) its sphere of influence and (b) its sphere of influence graph.

Computational geometry is another field where the term 'influence graph' has been previously used. Computational geometry concerns itself with designing and analyzing algorithms for solving geometric problems. In computational geometry, an influence graph is a generalized semi-dynamic structure which remembers the history of incremental construction of numerous geometric structures, for example, a Voronoi diagram, a structure used on a set of points (called *sites*) to solve proximity queries, such as the closest neighbors of a given site, or the closest pair of sites [12].

As can be seen, the new definition of the term 'influence graph' given in this paper is drastically different from the previously used definitions. This new definition of an influence graph may be useful in designing an intelligent agent to play the game of Diplomacy.

VI. PREVIOUS DIPLOMACY AGENT IMPLEMENTATIONS

The idea of creating an automated Diplomacy player first appeared some time in the early 1990s. It was quickly agreed upon that the most challenging task was deciding what units should move and where. In his 1994 article titled "Artificial Intelligence and Diplomacy", Steven Agar suggested that a library of openings be created, as it was done for chess, and quickly acknowledged that the moves following the opening moves would present a serious challenge [13]. He proposed two distinct approaches – considering factors to optimize in

turn the best move for each unit or considering the best moves for units to occupy specific positions on the board - and expressed his preference for the latter approach as the best way to formulate a coherent strategy that uses units together as opposed to all the moves being a collection of one-offs.

Games can be classified in several different ways: for example, as perfect or imperfect information or as games of chance. Diplomacy is not a game of chance, because no dice roll or anything else that randomly determines the outcome is involved, but it is a game of simultaneous moves where all orders of all players are kept secret until they are processed at the same time. This lack of knowledge of what the other players will do makes Diplomacy a game of imperfect information and greatly complicates matters.

In addition to Steven Agar's two approaches already mentioned – deciding on the best move for each unit or considering the best moves for units to occupy specific locations on the board – the third approach is a brute force exhaustive search of the move space, examining all possible combinations of moves by all players and picking the best move. The brute force approach relies on being able to search all possible moves (ideally over several turns) and to identify the result of each. The best move is the one with the least bad possible outcome, as this approach assumes perfect play by the opponents. Each leaf node of the game tree is assigned a score, with the path through the tree determined by both sides maximizing their own scores. Essentially any search trees paths possibly leading to defeat are avoided leaving only the paths that lead to a win or a draw against the perfect play [14]. Assigning a score to the leaf nodes of a game tree may be quite complex. The most apparent measure of a power's success in Diplomacy is the number of Supply Centers controlled by that power. However, many other considerations should be taken into account. For example, taking a supply center from an existing opponent may be more appealing than taking two supply centers from someone else. Other considerations may be the number and strength of the apparent enemies. A suitable score could be the difference between the number of Supply Centers that a power controls against a weighted average of everyone else, with the weights representing the degree of friendliness towards that power's rivals.

The advantage of the exhaustive search is that it provides for long term planning. The program implementing this search could find moves with no immediate advantage, but which will prove useful several turns later. This ability becomes crucial towards the end of the game where a power builds new units in its home centers, but it takes these new units several turns to reach the front lines. The advantage of the exhaustive search to its implementer is that the program would only need to know the rules of Diplomacy and have a good evaluation function. No other information would have to be provided. It would deduce tactics on its own.

While the exhaustive game tree search is a powerful method, it suffers from several major weaknesses. First of all, the search space is enormous with the exact number of unique openings being 4,430,690,040,914,420 not counting useless supports [15]. Second, in a multi-player environment, the strongest position both strategically and tactically is not necessarily the best position. For example, taking an early

lead in the game causes jealousy, and strong starters are likely to draw attention and get attacked from all sides. Finally, the exhaustive search ultimately fails, because Diplomacy is a game of imperfect information: all players move simultaneously, and no player knows what the other players are doing until all of their moves are processed at the same time. Finding the best moves requires knowledge of what the opponents' intentions are. Because the tree search assumes that the opponent will play perfectly and looks at the worst result of a set of moves, the risks necessary to advance may never be taken. This method lacks randomness. It finds the theoretically best move, but if this best move can be guessed by the opponent, it is clearly worse than a random move that still improves the overall position. But playing a less than perfect move against a perfect move by an opponent is not any better either. This is the weakness of the exhaustive tree search method in a game of imperfect information. The few existing Diplomacy implementations consider the current state of play and provide a set of moves aimed at improving the immediate situation by directing units to occupy certain locations. The longer term strategy is largely ignored.

An intelligent agent previously designed and implemented by Ritchie [16] scores every location on the map and moves the units to the highest scoring provinces. The factors considered when scoring locations are whether a province is a Supply Center (a more important factor in the Fall than in the Spring), who controls the province, whether it is threatened with an attack or whether it can be defended by an opponent and how important the neighboring locations are. To reduce predictability of a move, a small random element is added to each location. The size of this random element is worked out so that if two locations have exactly the same scores prior to the addition of the random element, the odds of choosing either location are exactly 50%. But even a small difference in scores of two provinces leads to a 60-40% or higher chance of the higher scoring provinces being selected for the move. Ritchie suggests an alternative solution where all provinces with a non-zero score would be selected according to the ratio of their scores. However, he did not add this solution to his implementation.

The location score is zero if the province is not threatened and it is 25 if an enemy threatens a Supply Center in that province. If an enemy Supply Center is threatened the province is worth 25 points or 40 points if the SC is undefended. But if the SC in the province being scored is an enemy home SC then the province is only worth 20 points or 60 if it is undefended, because it is inherently more valuable, and the enemy is more likely to defend it if it possibly can. 5 points are added to the score for each enemy unit adjacent to the province or 10 points if an enemy unit occupies the province. This encourages attacks against enemy units, cutting support that the enemy units can provide, and capturing enemy Supply Centers.

Each province receives a fraction of the adjacent provinces' scores to reflect the potential for the next season. This fraction depends on the season (one third in the Fall and one fifth in the Spring).

The next step in Ritchie's algorithm gives each unit a list of locations where it can move ordered by their score from highest to lowest, including the current location in case the

unit is already where it should be. Each unit attempts to move to the highest scoring location on its list. Before the move orders are finalized, the algorithm resolves conflicting moves so that if two units are attempting to move to the same location, one will either move to its second best location or simply support the move of the other unit.

Frederick Haard uses a similar tactical approach to unit moves, but his scoring is different. An agent is created for each unit of a power. An agent evaluates its surroundings and creates a goal list based on these evaluations. The goal list accounts for values and threats and includes data on how much support the unit requires for its moves to succeed as well as how much support it has available from the other units. A unit can have the goals of the other units on its list if it is expected to support them. Once a unit commits to a goal from its list, all its other goals are removed from its list as well as from the lists of the other units that may have supported these goals of the unit [17].

VII. EXAMPLES OF APPLYING INFLUENCE GRAPHS TO DIPLOMACY

As can be seen from the descriptions of the previous implementations of Diplomacy intelligent agents, their creators had an understanding that objects on the Diplomacy board – Supply Centers, navies and armies - have an influence on the value of the provinces they occupy as well as on the adjacent provinces. When scoring a province, Ritchie considers whether it is a Supply Center and whose unit occupies that province. The other factors, such as the control of the province, a threat to it, the ability of the opponent to defend it and the importance of the neighboring locations, are all reflections of the influence that various objects on the map have on provinces.

However, Ritchie's algorithm falls short of using the concept of influence graphs. The idea of influence graphs implies that the score of a vertex is influenced not only by the objects in the adjacent vertices, but by all the objects in the graph no matter what the distance is to the vertices occupied by these objects. Ritchie's algorithm does not consider the influence on a province by objects in the provinces further away than the adjacent provinces. For example, he adds 25 points to a province if it has a Supply Center and that SC is threatened by the enemy, meaning that an enemy unit is in a province adjacent to the province being scored. Adding points to the score in this situation is logical and justifiable as it makes the province more valuable and compels a power to defend it from the enemy threat. What is missing in this evaluation mechanism is the addition of points to a province's score if an enemy unit or units are within two or more *tempi* (plural of *tempo* – a term borrowed from chess, which in Diplomacy means a move of one piece from one province to another) of that province. This omission is partly what makes Ritchie's approach purely tactical and short sighted, designed to consider locations it currently identifies as important, as he himself acknowledged [18].

It will now be demonstrated what the author thinks influence graphs should look like in Diplomacy and how they should be combined to provide an evaluation of the Diplomacy board state. To make this demonstration easier, a

reduced Diplomacy map will be used. This map is a subset of the complete Diplomacy map, includes only Austria (we will shorten ‘Austria-Hungary’) and Germany, and excludes all maritime provinces (Figure 11). It uses only one type of units – land armies. The reason for selecting this submap is that these two powers are comprised of six provinces each, have the same number of home Supply Centers (three), and have an extended border, which allows to avoid stalemates early in the game after only a couple of moves. It is also best to use a reduced version of the game with only two players as it leaves out the diplomatic negotiations, which are an essential part of Diplomacy, but are outside the scope of this paper.

Provinces on this map are influenced by two types of objects – Supply Centers and armies belonging to either power. Since Supply Centers are in fixed locations, an influence graph for a Supply Center will be generated first.



Figure 11: Reduced Diplomacy map

A Supply Center exerts the greatest influence on the province which it occupies. It exerts some influence on all other provinces on the map, but its influence on a province wanes with distance between the Supply Center and that province. Figure 12 shows the normalized map as a graph where provinces are vertices. All provinces have their names abbreviated to just three letters. If two provinces have a common border the corresponding vertices are connected by an edge.

Of all six Supply Centers, Vienna (VIE) occupies the most centralized location (all other Supply Centers are located in provinces bordering on an edge of the map). So its influence graph will be constructed first as this graph will best illustrate the concept. An arbitrary score of fifty will be assigned to the province where Vienna is located due to the influence of the Supply Center. Then the graph where vertices represent provinces is traversed and the shortest path from Vienna to each of the other provinces on the map is identified. The shortest paths of length one will be from Vienna to Tyrolia, Bohemia, Galicia, Budapest, and Trieste. The longest paths of length three will be from Vienna to Ruhr, Kiel, Berlin, and Prussia.

Depending on the length of the path, the influence of the Supply Center in Vienna will decrease by a certain degree. For a path of length one, Vienna’s influence will decrease by 20%. For a path of length two, it will decrease by 40% and so on with each additional path length unit decreasing Vienna’s

influence by an additional 20%. Figure 13 shows Vienna’s influence graph.

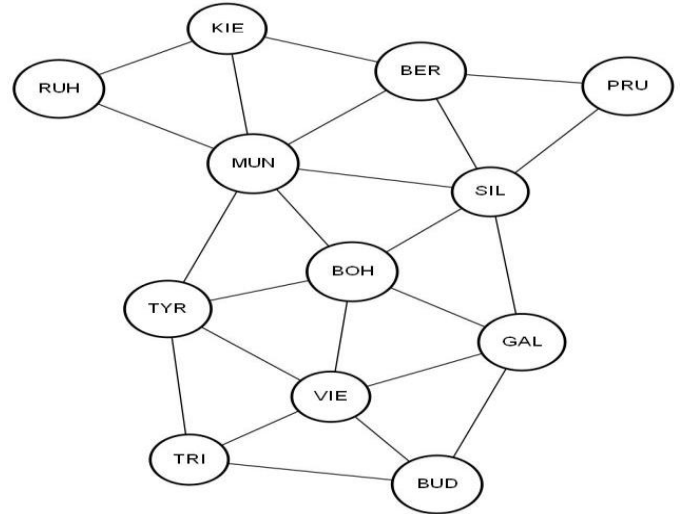


Figure 12: Reduced Diplomacy map as a graph

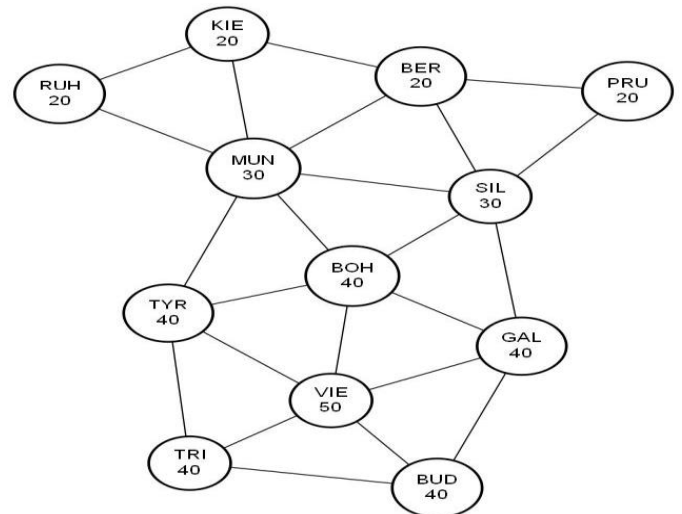


Figure 13: Vienna’s influence graph

Because Supply Centers are in fixed locations, their influence does not change in the course of the game. The armies, on the other hand, – both German and Austrian – move in the course of the game. The goal of Diplomacy is to capture as many enemy Supply Centers as possible while not losing one’s own centers to the enemy. Trying various sets of weights for Supply Centers shows that enemy Supply Centers should have bigger weights than home Supply Centers in order to have a board evaluation closer to that of a human player. This distinction calls for creating two models of the world – one for each power. Vienna’s influence graph in Figure 13 is to be used when creating Germany’s model of the world: to Germany Vienna is an enemy Supply Center which is to be captured. To Austria Vienna is a home Supply Center which is to be defended. Vienna will be assigned an influence weight of ten in its province in Austria’s ‘world’ (Figure 14).

Another reason for generating a separate and distinct ‘world’ for each of the two powers is that their armies should have an opposite influence on each other. Austrian units have a positive influence on provinces in Austria’s world, but they

exert a negative influence on the value of provinces from the German point of view. Several trials of different weights have shown that an acceptable absolute value for an army is 15.

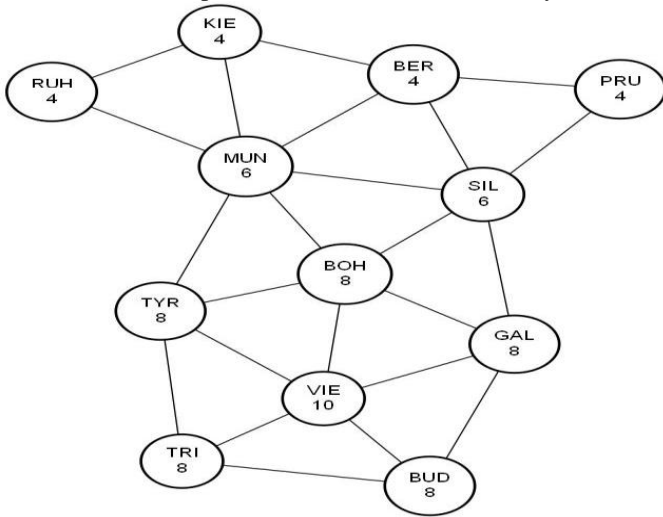


Figure 14: Vienna's influence graph (in Austria's world)

Figure 15 shows the influence graph for an Austrian army unit initially located in the province of Vienna. It is this army's influence from Austria's perspective. From Germany's perspective, the influence of this Austrian unit will be exactly opposite (negative).

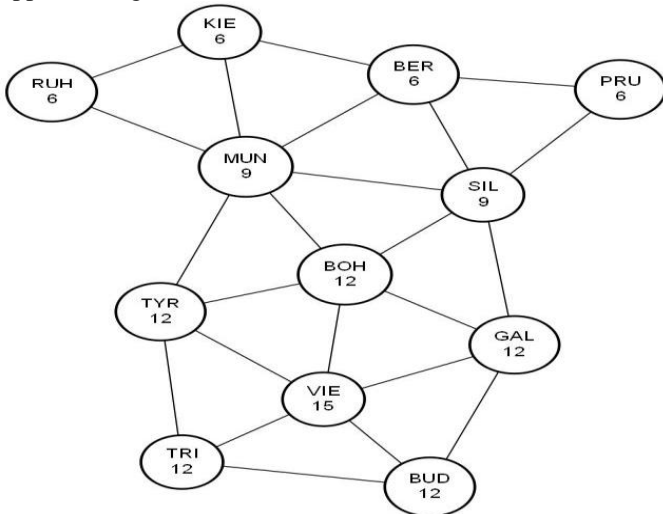


Figure 15: Influence graph of the army at Vienna

At the start of the game there are a total of twelve influence graphs for each of the two worlds – one for each Supply Center and each army. The graphs are combined by adding the influence on each province (cell) by each Supply Center and each army. For example, the score of Vienna in the resulting Austrian influence graph would include ten points because a Supply Center is located in this province (Figure 14), fifteen points for the Austrian army unit located there (Figure 15) and additional points due to the influence of the other five Supply Centers and five armies on the board.

Let us construct two combined influence graphs for the initial state of the game. At the start of the game the three armies of each power occupy that power's Supply Centers: Germany has units placed at Kiel, Munich, and Berlin.

Austria has units placed at Vienna, Trieste, and Budapest. Figure 16 shows the initial state of the game.

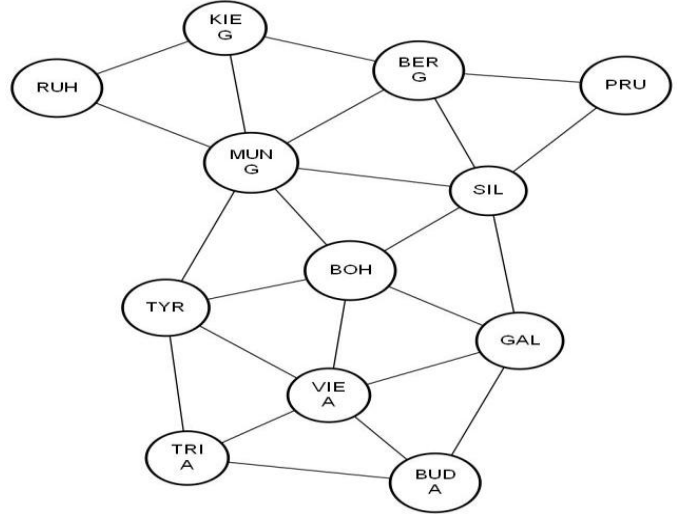


Figure 16: Initial state of the game (A – Austrian army, G – German army)

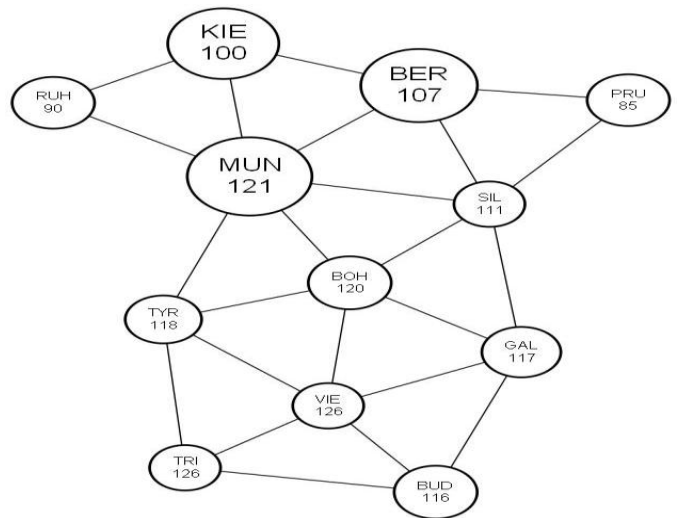


Figure 17: Combined German IG at the start of the game

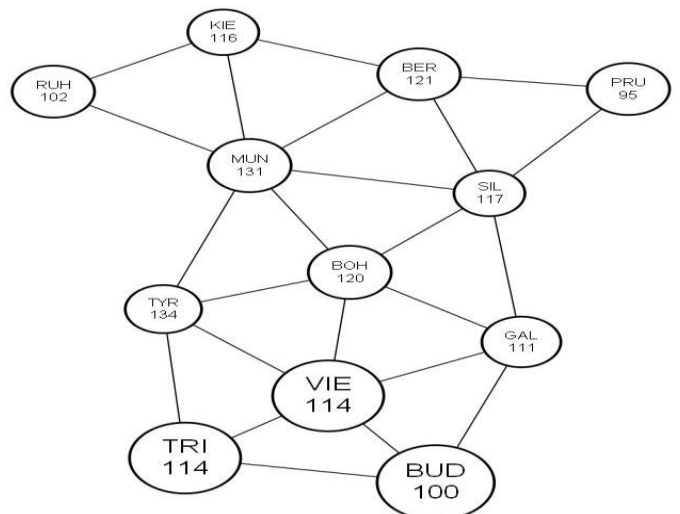


Figure 18: Combined Austrian IG at the start of the game

At the start of the game the German combined influence graph is shown in Figure 17 and the Austrian combined influence graph in Figure 18.

Austria occupies three provinces with a combined score of 328 out of the total score of 1375 for all twelve provinces. Thus, Austria's standing is $(328/1375) = 23.85\%$. The German armies occupy three provinces with the same combined score of 328, but from Germany's point of view all twelve provinces' total score is only 1337. Germany's standing is $(328/1337) = 24.52\%$. Thus, Germany's position at the start of the game appears slightly superior compared to Austria's.

Let us suppose that the German armies (circles) have somehow managed to advance deep into the Austrian territory and to push the Austrian armies (squares) back to their original positions. The corresponding state of the game is shown in Figure 19. The combined influence graphs for this state of the game are shown in Figures 20 and 21.

The combined score of the provinces Germany occupies is 382 out of the total score of 1355 or 28.19% while the combined score of the provinces occupied by Austria is 289 out of the total score of 1357 or only 21.30%. Thus, Germany's position is superior to Austria's which would match the assessment of an experienced Diplomacy player: the German units are in the provinces adjacent to the Austrian Supply Centers and are threatening these centers. The German Supply Centers are under no Austrian threat whatsoever.



Figure 19: Germany's advantage over Austria

Let us now look at an opposite scenario: the Austrian units (squares) have either advanced into German territory or are about to enter it and have pushed the German units (circles) back to their original positions. The corresponding state of the game is shown in Figure 22.

The combination of influence graphs for this state of the game is shown in Figures 23 and 24. The combined score of the provinces Austria occupies is 386 out of a total score of 1420 (27.18%). The combined score of the provinces Germany occupies is 292 out of a total score of 1292 (22.60%). Thus, Austria's position is superior, which would yet again match the assessment of an experienced Diplomacy player: the Austrian units are in the provinces adjacent to two German Supply Centers and are threatening one of these centers while the Austrian Supply Centers are under no German threat whatsoever.




Figure 20: German IG for the game state in Figure 19

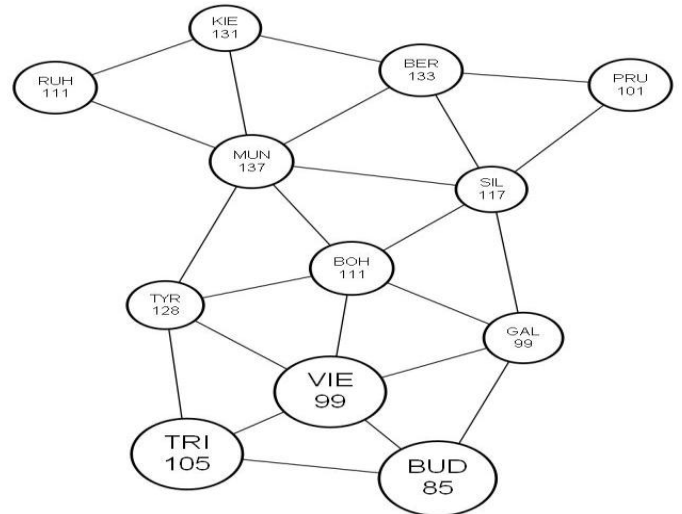


Figure 21: Austrian IG for the game state in Figure 19

The full Diplomacy board was designed in a way which would make all seven powers approximately equal in strength at the start of the game. Even though only an arbitrary subset of the Diplomacy board is used in this paper and the influence of the other five powers on Austria and Germany is disregarded, Germany and Austria are still of approximately equal strength, according to their combined influence graphs. This could be due to a 'good' selection of a submap, which preserves a close balance of power between two countries even in the absence of the other five powers. In any case, so far the use of influence graphs confirms the expectation that two powers should be approximately equal at the start of the game. Combined influence graphs will now be constructed for some 'extreme' board states that appear to favor one of the powers over the other. These states may or may not actually arise in the course of the game, but they should give us a sense of accuracy in evaluating the state of the board that influence graphs provide.



Figure 22: Austria's advantage over Germany




Figure 24: Austrian IG for the game state in Figure 22

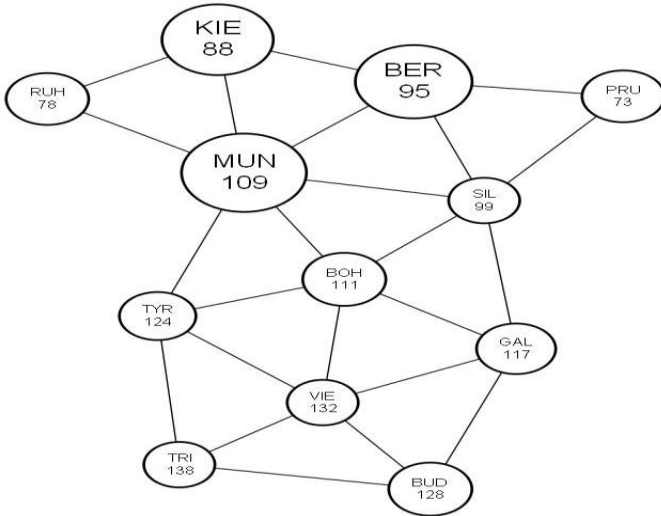


Figure 23: German IG for the game state in Figure 22

The longest shortest path between two provinces on the France-Germany map (Brest-Prussia) is five while the longest shortest path on the Austria-Germany map was four. This path of length five is the only one in the graph. The influence of the Supply Center and any unit that may be at Brest will decrease at the same rate along paths to other provinces as the influence of any other Supply Center and any army on the Austria-Germany and France-Germany maps. The influence of Brest and any army that may be at Brest on Prussia will be zero. The influence of any unit which may be in Prussia on the province of Brest will be zero as well. The combined influence graphs for this game state show that the German armies occupy provinces with a combined score of 377 out of a total score of 1318 (28.60%) while the French armies occupy provinces with a combined score of 265 out of a total score of 1262 (or 21%).

Even though in both preceding examples influence graphs provide an evaluation of the game state matching that of an experienced human Diplomacy player, it cannot be ruled out that this coincidence is purely accidental. After all, the Austria-Germany map is an arbitrary subset of the full Diplomacy map, which breaks down multiple 'checks and balances' other five powers have in relation to Germany and Austria. In order to demonstrate that the board evaluation method using influence graphs is accurate, a different Diplomacy submap will be considered, which includes only France and Germany.

First, the state of the game will be evaluated at a point where the German armies (circles) have somehow managed to advance into the French territory while the French armies (squares) have been pushed back to their initial positions. The corresponding state of the game is shown in Figure 25. Clearly, from a human player's point of view Germany has an advantage over France at this point in the game. Let us now see whether a combined influence graph supports this assessment.



Figure 25: Germany's advantage over France

In an opposite game state, a French army (square) has advanced into the German territory and two other French armies (squares) are either on the border or close to it while the German armies (circles) are in their initial positions (Figure 26). France clearly has an advantage over Germany in this state of the game, even though France may be unable to advance any further. The combined influence graphs for this game state show that the German armies occupy provinces with a combined score of 291 out of a total of 1219 (23.87%) while the French units occupy provinces with a combined

score of 367 out of a total of 1361 (or 26.97%). In both cases shown in Figures 25 and 26 the combined influence graphs provide a game state evaluation which favors the same power as an experienced human Diplomacy player would.

It was stated earlier that building and using a game tree for exhaustive search is not practical in Diplomacy due to the immense size of the tree. Yet, analyzing parts of the search tree using a submap may be helpful in demonstrating the utility of influence graphs. It will now be shown how influence graphs can be used in a game tree search.

VIII. INFLUENCE GRAPHS AS BOARD STATE EVALUATORS IN GAME TREE SEARCH

To decide what the best move for an army is, an intelligent agent must take into account various factors from the army’s surroundings, such as the number and proximity of enemy armies, the number and proximity of its own or allied armies which it can either support or receive support from, the proximity of own or allied Supply Centers and whether these centers are threatened and need to be defended, the proximity of enemy Supply Centers and the desirability of attacking and capturing them, etc. Because exhaustive search is not practical due to the enormous size of the search tree, we will use a different kind of search for a Germany-Austria subgame to demonstrate the utility of influence graphs in evaluating the state of the game.



Figure 26: France’s advantage over Germany

When it is impractical to construct a complete search tree, search should be cut off earlier and a *heuristic evaluation function* should be applied to states in the search, effectively turning nonterminal nodes into terminal nodes. A utility function used in exhaustive minimax search is replaced by a heuristic evaluation function, which gives an estimate of the position’s utility and the terminal test is replaced by a cutoff test which decides when to apply the evaluation function. An evaluation function returns an estimate of the expected utility of the game from a given position. The evaluation function should strongly correlate with the actual chances of winning [19]. The rest of this section demonstrates how influence graphs can be used in creating a heuristic evaluation function for a Diplomacy tree search.

There are multiple sets of opening moves for Austria and Germany even if useless ‘stands’ (where an army just stands instead of moving) and circular ‘swaps’ (all three armies belonging to the same power moving in a circular fashion) are

excluded. An example of an excluded set of moves with a stand is (BER=>SIL, MUN=>BOH, KIE stands). Some opening moves would seem ‘irrational’, such as BER moving to PRU or KIE moving to RUH (see Figures 11 and 12 for the correspondence of full and abbreviated province names): they are moving German armies away from the enemy. The weakness of these moves is obvious to a human Diplomacy player and is confirmed by the influence graph based game state evaluator. Most sets of opening moves which include such ‘irrational’ moves are excluded from consideration due to space constraints, but two such ‘irrational’ German moves will be considered to illustrate their ineffectiveness. Also excluded are sets of moves which lead to the same positioning as does another set which has already been considered. For example, a set of Austrian moves (TRI=>TYR; VIE supports TRI, BUD=>TRI) leads to exactly the same positioning of the Austrian armies as does another set (VIE=>TYR; TRI supports VIE; BUD=>VIE) and is therefore excluded from consideration. There are eight sets of opening moves by Germany to be considered:

G1	G2	G3
BER => SIL MUN => TYR KIE => MUN	BER => SIL MUN => TYR KIE => BER	BER => SIL MUN => BOH KIE => MUN
G4	G5	G6
BER => SIL MUN => BOH KIE => BER	BER => SIL MUN => BOH KIE => RUH	BER => SIL MUN => TYR KIE => RUH
G7	G8	
BER => PRU MUN => TYR KIE => BER	BER => PRU MUN => TYR KIE => MUN	

Table 1: German opening moves

The Austria-Germany map is structured in such a way that there are no ‘irrational’, retreating opening moves for Austria. This power has a choice of seven opening moves that should be examined:

A1	A2	A3	A4
VIE => BOH TRI => TYR BUD => VIE	VIE => BOH TRI => TYR BUD => TRI	VIE => BOH TRI => TYR BUD => GAL	VIE => TYR TRI supp VIE BUD => VIE
A5	A6	A7	
VIE => TYR TRI supp VIE BUD => GAL	TRI => TYR VIE supp TRI BUD => GAL	VIE => BOH TRI => VIE BUD => GAL	

Table 2: Austrian opening moves

In game theory simultaneous move games like Diplomacy can be presented in *extensive form* as a tree or in *normal form* as a payoff matrix [20]. In the extensive form of the Austria-Germany game, the root of the tree will represent the first move made by Germany. The root will have eight branches to nodes representing the four possible sets of German opening moves. Each of these eight nodes will have seven branches to the same seven nodes representing seven sets of possible Austrian opening moves or reactions to the German opening moves. Because Diplomacy is a simultaneous move game, the top two levels of the tree represent one ply (Figure 27).

The purpose of searching the tree is to find a leaf representing a win (in reality, because we are not running an exhaustive search, we are looking for a node that is *likely* to lead to a win if it is expanded). A win is defined as capturing at least one enemy supply center during a Fall season or occupying it in a Spring season and controlling it by the following Fall season. Due to the small size of the Germany-Austria submap, it is very likely that some moves will lead to a stalemate. Some leaves will be stalemates. A stalemate is defined as a game state where a power, if the rival were to stand, can only make moves which the influence graph-based game state evaluator shows to be weakening that power's position on the board. In this experiment, it is desirable to find either a win by one of the powers or a stalemate.

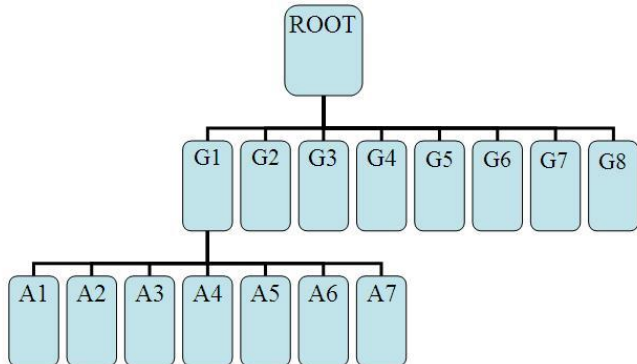


Figure 27: Partial view of the Diplomacy search tree

In this experiment the same weights are used as in the examples in part VII: home Supply Centers are given a weight of ten each while enemy supply centers are given a weight of fifty each. Own armies have a weight of +15 each while the enemy armies have a weight of negative 15 each. The position of a power is measured, as before, by the percentage of the combined score of the provinces occupied by its armies in relation to the combined score of all twelve provinces in that power's 'world'. The example in Figure 18 shows the Austrian 'world' at the start of the game. While the total score of all provinces for Austria is 1375, the three Austrian armies occupy three provinces (marked by a large font) with a combined score of 328. Thus, the current standing of Austria in the game is rated at $(328/1375) = 23.85\%$.

To evaluate the state of the game at a node in the search tree, the heuristic evaluation function needs to measure Austria's standing against Germany's standing. The author proposes the ratio (expressed as a percentage) of Austria's standing to Germany's standing as such measurement. For example, at the start of the game, Austria's standing is at 23.85% (Figure 18) while Germany's standing is at 24.52% (Figure 17). The ratio of their standings is $(23.85\%/24.52\%) = 97.24\%$. A value under 100% would indicate that Germany has an advantage. A value over 100% would indicate an Austrian advantage. The payoff matrix for ply one search tree nodes is presented in Table 3.

The game theory tells us that at this point in this simultaneous game each power, in order to maximize its payoff, must either choose a dominant pure strategy or, in the absence of a dominant strategy, use a mixed strategy, which is a combination of several pure strategies, one of which is

picked at random with a certain probability that can be calculated [21]. One of a player's pure strategies is said to *dominate* another strategy if it yields an outcome at least as good against any of the pure strategies that his opponent may choose [22]. If each power has a dominant strategy, the point in the payoff matrix where these two strategies intersect is called a *saddle-point* [23]. Let us try to determine if one or both powers have a dominant strategy.

Strat	A1	A2	A3	A4	A5	A6	A7
G1	103.05	97.41	102.68	108.96	106.81	105.54	92.15
G2	99.41	94.01	99.04	105.16	103.10	101.82	97.67
G3	100.25	108.96	105.54	97.16	95.06	93.75	98.92
G4	96.78	105.16	101.82	102.77	100.69	99.45	95.50
G5	99.72	108.41	105.03	105.84	103.72	102.57	98.41
G6	101.06	95.58	102.17	108.41	106.27	105.03	100.64
G7	107.85	102.09	107.67	113.83	111.88	110.64	105.95
G8	112.33	106.15	112.06	118.50	116.34	115.12	101.35

Table 3: Ply one of the search tree in normal form

The search can begin by *inadmissible* strategies being eliminated. A pure strategy which is dominated by another strategy is inadmissible in the terminology of the game theory, and a rational player will never choose it [22]. Austrian strategy A7 is dominated by strategy A3 and strategies A5 and A6 are dominated by strategy A4. Therefore, columns representing strategies A5, A6, and A7 can be deleted from the payoff matrix, since a reasonable Austrian player would never use them. What is now left is a 7x4 matrix, in which the rows representing German pure strategies G1, G5, G6, G7, and G8 are dominated by the row representing German pure strategy G2. Thus, rows G1, G5, G6, G7, and G8 can also be eliminated from the payoff matrix. After dominated rows and a dominated column are eliminated, a 3x4 payoff matrix is left in Table 4:

Strategy	A1	A2	A3	A4
G2	99.41	94.01	99.04	105.16
G3	100.25	108.96	105.54	97.16
G4	96.78	105.16	101.82	102.77

Table 4: Payoff matrix after dominated rows and columns are eliminated.

No more pure strategies can be eliminated after this. The game theory tells us that a mixed strategy is then to be used. However, it can be shown, by expanding the three nodes which include A3 that this pure strategy leads to an Austrian victory. This leads the author to believe that A3 is the dominant opening moves strategy, but that the current data does not support it because the selected influence weights of Supply Centers and armies and the rates of their influence decrease are not good enough to capture all the nuances of the game. Further experimentation with these constants is required, probably using genetic algorithms, to optimize the weight and rate of influence decrease constants. If Germany had a dominant strategy as a result of constants fine tuning, for example G2, then the tree node G2A3 would have to be expanded. If Germany still did not have a dominant strategy then, according to the game theory, it would have to use a mixed strategy by choosing one strategy from G2, G3, and G4 at random. The odds of choosing one of these three strategies

can be calculated, but these calculations will be left out and, instead, it will be shown what happens in the case of Germany choosing the first one of these three strategies, G2. In other words, the tree node G2A3 will be expanded.

Tree node G2A3 represents the state of the game after the moves in strategies G2 and A3 have been processed and all conflicts have been resolved. This node will have four branches for four possible German sets of moves (Table 5). Each of the nodes at the end of these four branches will have three branches for three possible Austrian moves in response (Table 6).

G1	G2	G3	G4
SIL => BOH	MUN => BOH	SIL => GAL	SIL => GAL
MUN supp SIL	SIL supp MUN	MUN => TYR	MUN => TYR
BER => SIL	BER => MUN	BER => SIL	BER => MUN

Table 5: Possible sets of German moves from node G2A3

A1	A2	A3
TRI => TYR	TRI => TYR	BOH => SIL
BOH supp TRI	GAL => SIL	GAL supp BOH
GAL => SIL	BOH supp GAL	TRI => TYR

Table 6: Possible sets of Austrian moves from node G2A3

Node G2A3 of the first ply is thus expanded into twelve nodes in ply two. This part of the search subtree is presented in its normal form in Table 7:

Strategy	G1	G2	G3	G4
A1	98.69%	107.31%	107.31%	107.31%
A2	98.69%	107.31%	112.99%	112.99%
A3	107.31%	109.49%	108.47%	108.47%

Table 7: Result of expanding node G2A3

The payoff matrix in Table 7 has a *saddle-point*. It is situated at the intersection of row 3 and column 1: 107.31% is a minimum in its row and a maximum in its column. The existence of the saddle-point means that Germany should choose strategy G1 and Austria should choose strategy A3. If they both are reasonable and choose these strategies the resulting game state is pictured in Figure 28:

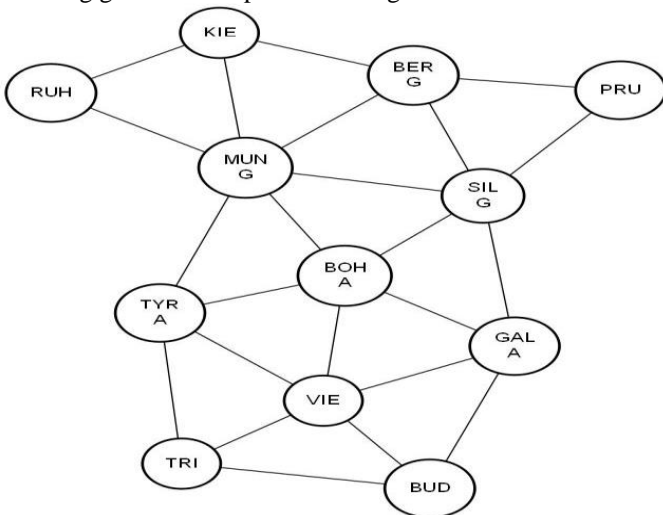


Figure 28: Game state after both player select dominant strategies when ply one node G2A3 is expanded (A – Austrian army, G – German army)

If Germany chooses a different strategy as its set of opening moves – either G3 or G4 – expanding nodes G3A3 and G4A3 produces payoff matrices with dominant strategies which lead, just like the expansion of node G2A3 did, to the game state pictured in Figure 28. A further expansion of this state (ply three of the search tree) will lead to Austrian army at GAL capturing SIL with support from the army at BOH while the German army at SIL will have to retreat to PRU, the only province available for retreat. The German army at BER will not be able to provide support to the army at SIL: it will have to support the army at MUN to prevent a successful attack on MUN by two Austrian armies at BOH and TYR which would lead to a German loss of MUN and, thus, of the game. After that (ply four of the search tree) Austria’s unit at SIL is positioned to capture the German supply center at MUN with support from the armies at BOH and TYR thus winning the game for Austria (capturing an enemy supply center was defined as victory and is therefore a leaf or terminal node of the search tree).

IX. OTHER APPLICATIONS OF INFLUENCE GRAPHS

Influence graphs are applicable to games with boards that can be represented as a general graph, but not as a grid, such as *Risk* or any other board game that uses either a real or a fictional geographical map.

Influence graphs could be useful in some real world applications, which use geographical maps. In 2006 the Russian State Duma (parliament), in response to public concerns over the negative social consequences of gambling, passed a law which would ban gaming establishments anywhere in the country by July 2009 except for four designated ‘gaming zones’. The four zones have been identified as Kaliningrad (West), Primorie (Far East), Altai (Siberia) and an area on the border of Krasnodar and Rostov regions (South) shown in Figure 29.



Figure 29: Russia’s four proposed gaming zones.

The first three of the four future gaming zones cover entire administrative regions of the country. There has been little

explanation how these four zones were selected, but it is possible to deduce some reasons behind this selection by keeping in mind the aim of this legislation and by analyzing the geographical location of the four proposed zones. The goal of the legislation was to relegate gambling to remote areas of the country away from large population centers and yet to allow the gambling industry to operate in the areas where it could take maximum advantage of existing infrastructure and draw the maximum number of domestic and foreign visitors. The use of influence graphs with each accounting for a factor, such as the distance of all of the country's regions from large population centers, their proximity to major transportation routes, their level of infrastructure development, proximity to foreign countries, availability of local workforce, construction costs and others, could help designate the optimal regions as gaming zones.

X. CONCLUSION AND FUTURE WORK

The game of Diplomacy provides a good example of a world model which cannot be represented or treated as a grid. This makes influence maps inapplicable to Diplomacy and other games with a board best represented by a graph. Influence graphs which were defined and whose use was demonstrated in this paper are a helpful tool in evaluating the state of the Diplomacy 'world' and may be useful in evaluating the state of other games and of the world in real life applications.

Much can be done to further develop and improve the influence graphs even in their application to Diplomacy submaps used in this paper. The constants for the weights of the Supply Centers and the armies and the rate of the decrease of their influence with distance have been selected arbitrarily and then slightly tuned so that the world evaluation in which they are used only roughly matches the judgment of an experienced human Diplomacy player. Genetic algorithms could be used to optimize these constants, although a good fitness function would present a tremendous challenge.

REFERENCES

- [1] Russell, S. & Norvig, P. (2003). *Artificial Intelligence: A Modern Approach* (2nd ed.) Upper Saddle River, New Jersey: Prentice Hall. p. 161.
- [2] Miles, C & Louis, S. (2006). *Co-evolving real-time strategy game playing influence map trees with genetic algorithms*. Retrieved January 19, 2009 from <http://www.cse.unr.edu/~sushil/pubs/newestPapers/2006/cec2006/cec2006.pdf>
- [3] Sweetser, P. & Wiles, J. (2005). *Combining Influence Maps and Cellular Automata for Reactive Game Agents*. Retrieved November 24, 2008 from http://www.emergenceingames.com/web_images/Sweetser_I_DEAL05.pdf
- [4] Sharp, R. (1978). *The game of Diplomacy*. London: Arthur Barker. p. 1
- [5] Ritchie, A. (2003). *Diplomacy AI*. Unpublished Master Thesis. University of Glasgow, Scotland. pp. 5-6
- [6] *Webster's college dictionary* (2001). Random House. p. 540
- [7] Crockford, D. *Crockford Diplomacy map*. Retrieved April 21, 2009 from <http://www.crockford.com/wrrld/diplomacy.html>
- [8] Chachra, V., Ghare, P. M., & Moore, J. M. (1979). *Applications of graph theory algorithms*. Elsevier North Holland. pp. 54-57
- [9] Shortest Paths Algorithms: Dijkstra's Greedy Algorithm. Retrieved May 4, 2009 from http://www.cacr.caltech.edu/~sean/projects/stlib/html/shortest_paths/shortest_paths_dijkstra.html
- [10] Chalker, T.K., Godbole, A.P., Hitczenko, P., Radcliff, J. & Ruehr, O.G. (1998). *On the size of a random sphere of influence graph*. Retrieved March 9, 2009 from <http://www.etsu.edu/math/godbole/SIGFINAL.pdf>
- [11] Soss, M.A. (1998). *On the size of sphere of influence graph*. Retrieved March 9, 2009 from http://digitool.library.mcgill.ca:8881/dtl_publish/6/20869.html
- [12] Teillaud, M. (1993). *Towards dynamic randomized algorithms in computational geometry*. Springer-Verlag.
- [13] Agar, S. (1994). *Artificial Intelligence and Diplomacy*. Retrieved January 20, 2009 from <http://www.diplomacy-archive.com/resources/postal/artificial.htm>
- [14] Ritchie, A. (2003). *Diplomacy AI*. Unpublished Master Thesis. University of Glasgow, Scotland. p. 18
- [15] Haard, F. (2004). *Multi-Agent Diplomacy. Tactical planning using cooperative distributed problem solving*. Unpublished Master Thesis. Blekinge Institute of Technology, Sweden. p. 1
- [16] Ritchie, A. (2003). *Diplomacy AI*. Unpublished Master Thesis. University of Glasgow, Scotland. pp. 20-21
- [17] Haard, F. (2004). *Multi-Agent Diplomacy. Tactical planning using cooperative distributed problem solving*. Unpublished Master Thesis. Blekinge Institute of Technology, Sweden. pp. 12-16
- [18] Ritchie, A. (2003). *Diplomacy AI*. Unpublished Master Thesis. University of Glasgow, Scotland. p. 21
- [19] Russell, S. & Norvig, P. (2003). *Artificial Intelligence: A Modern Approach* (2nd ed.) Upper Saddle River, New Jersey: Prentice Hall. pp. 171-173

[20] Colman, A. (1982). *Game theory and experimental games*. Pergamon Press. pp. 49-50

[21] Colman, A. (1982). *Game theory and experimental games*. Pergamon Press. pp. 51-68

[22] Colman, A. (1982). *Game theory and experimental games*. Pergamon Press. p. 61

[23] Colman, A. (1982). *Game theory and experimental games*. Pergamon Press. p. 52