A Satis£cing Approach to Assigning Vehicles to Targets *

Michael A. Goodrich

Computer Science Department
Brigham Young University
Provo, UT, USA 84602

Abstract

Determining which of N-agents are assigned to which of M-tasks is order M^N in complexity. By requiring individual agents to identify the set of individually satis£cing actions, this complexity remains exponential but can be signi£cantly reduced. We present an algorithm for identifying individually satis£cing choices in a problem where agents must reach a target while avoiding threats. We demonstrate that the algorithm provides some computational savings and produces reasonable assignments.

1 Introduction

Multi-agent and multi-robot teams are becoming increasingly important in the research community. Often, it is desirable to assign these robots to particular tasks so that multiple tasks are accomplished in a distributed manner. Unfortunately, optimally assigning agents to tasks quickly becomes computationally intractable as the number of agents increases. In the absence of an optimal solution, it is necessary to £nd algorithms that produce solutions that are good enough, or *satis£cing* [9], in a tractable number of computations.

In this paper, we present a technique for producing such solutions. Our approach to the assignment problem is based on a simple principle: the set of socially optimal allocations should be constrained by the set of individually satis£cing assignments. In other words, in a problem where each agent must perform a task, no agent should be assigned to a task which involves too much risk or work.

After presenting a brief review of related literature, we pose the problem of assigning agents to targets subject to risks. We then evaluate the computational requirements of our proposed solution, and characterize typical assignments for a simulated environment.

2 Related Literature

The assignment problem is a well known optimization problem. In its most simple form, the problem is to assign items (jobs, missiles) to other items (machines, targets) [4]. Such problems frequently arise in the literature, especially in multi-agent settings such as task-decomposition problems [11] and auction-based task assignment [3]. Although there are algorithms for solving the assignment problem subject to a performance criterion, the problem is NP-hard whence heuristic techniques are often used [2, 7].

In our treatment of the target assignment problem, we use satis£cing decision theory [5] to balance the desires of individuals with the needs of the team. This theory, which has been applied in both control settings [6] as well as in multi-agent interactions [10], has been applied to formation-initialization problems wherein agents (robots, satellites) are assigned to a location in the formation [1].

3 The Problem

Consider a world like the one illustrated in Figure 1. In the world, there are M targets denoted by the bullseye symbols, N vehicles denoted by the triangle symbols, and P threats denoted by the T symbols. The task is to assign each vehicle to a target in a sensible way that avoids threats, avoids long paths, and covers a lot of targets.

*Accepted for publication in 2001 International Conference on Systems, Man, and Cybernetics. Due to a publishing error, this paper did not appear in the proceedings.

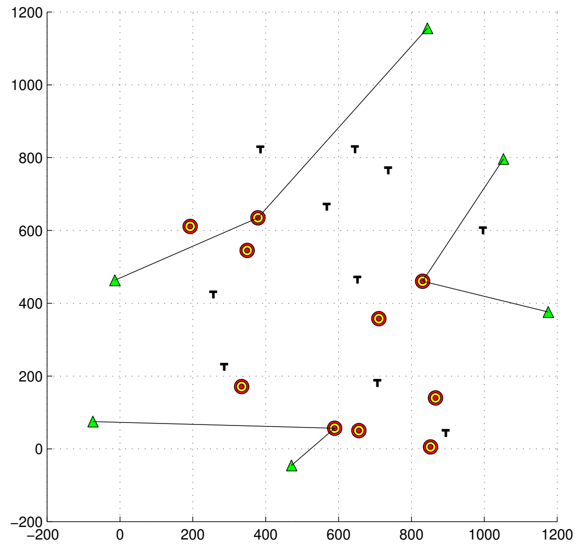

Figure 1: A world in which N vehicles must choose from M targets and void P threats. Included in the £gure is a typical assignment, indicated by a solid line and generated by our algorithm, of targets to vehicles. A solid line connects each vehicle to its assigned target. Notice how balanced teams naturally form if threats can be avoided.

3.1 Complexity of Optimal Solution

The assignment of agents to tasks is inherently a centralized task in that the resulting solution must produce some socially useful distribution. (Even solutions that are determined via negotiations between distributed agents must be in¤uenced by a negotiation protocol that is designed to produce a socially desirable solution.) For decision problems such as the one outlined in Figure 1, the standard solution approach is to create an objective function that encodes all of the desiderata of the decision problem. The optimal solution is one that maximizes this objective function. We can £nd the expected complexity for £nding an optimal solution to this target assignment problem. For N agents with M targets, we must search through $\prod_{i=1}^{N} M = M^N$ possible assignments. For the world diagrammed in Figure 1, $N = 6$ and $M = 10$ whence the time complexity of £nding an optimal solution is $M^N = 1,000,000$. Clearly, an exhaustive search through such a world is infeasible for large numbers of vehicles and many targets.

3.2 Measures of Goodness

A brief analysis of the problem alerts us to four objectives:

- minimize the distance between the location of a vehicle's starting point and the location of its assigned target (that is, we want to minimize the distance that each vehicle must travel),
- maximize the distance between all points on the vehicle's path to a target and each threat,
- maximize the number of vehicles assigned to each target (to maximize ef£ciency), and
- maximize the number of targets attacked.

We refer to these objectives as the *ShortPath*, *AvoidThreats*, *MaxForce*, and *MaxSpread* heuristics, respectively.

4 Individually Satis£cing Choices

Of the four objectives, the *ShortPath* and *AvoidThreats* heuristics are based on the desires of individual agents whereas the *MaxForce* and *MaxSpread* heuristics are based on the desires of the team. We propose a solution methodology that restricts the set of feasible team assignments to those that are acceptable to individual agents. In other words, we let individual agents eliminate those targets that expose them to too much risk or that require too great a distance to travel, and then maximize team choices within this individually satis£cing set. We will demonstrate a dramatic decrease in time complexity for this approach.

4.1 *ShortPath* Heuristic

From an individual perspective, each vehicle must accept an assignment to attack one and only one target. From this myopic perspective, targets that are closer to a vehicle are preferred to targets that are farther away. For each potential assignment for vehicle V_i to target G_j, we can denote the straight-line (Euclidean) distance between the vehicle and the target by $d_A(V_i, G_j)$.

The measure of *acceptability* of the assignment of vehicle V_i to target G_j grows as the distance $d_A(V_i, G_j)$ decreases. If we denote this measure of acceptability by $\mu_{A_i}(G_j)$, we can relate μ_{A_i} and $d_A(V_i, G_j)$ as follows:

$$\mu_{A_i}(G_j) = \frac{\max_{G_j}\Big(d_A(V_i,G_j)\Big) - \Big(d_A(V_i,G_j)\Big)}{\max_{G_j}\Big(d_A(V_i,G_j)\Big) - \min_{G_j}\Big(d_A(V_i,G_j)\Big)}.$$

This transformation gives the closest target the highest measure of acceptability (equal to one), the most distant target the lowest measure of acceptability (equal to zero), and all other possible targets an assignment in the range $[0, 1]$. Notice that this function can be computed and searched in order NM time.

4.2 *AvoidThreats* Heuristic

From an individual perspective, each individual wants to avoid threats. We can construct a measure of the *rejectability* of each assignment by estimating how close each agent will get to a threat on its path to the target. We denote this measure of rejectability by μ_R, and construct it in two steps. We will refer to Figure 2 throughout the discussion, and we will let $d_R(V_i, G_j, T_k)$ denote the threat caused by T_k to vehicle V_i if V_i is assigned target G_j.

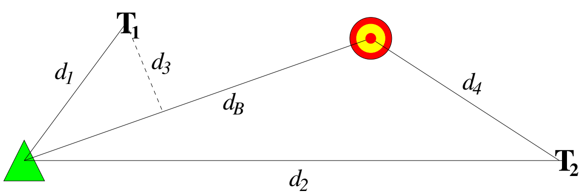

Figure 2: Computing how close a vehicle comes to a threat while traveling to a target.

The £rst step is to estimate how much threat exposure each agent will experience on its path to each target. Consider the relationship between the vehicle and T_2 in Figure 2. In the £gure, the distance d_B between the vehicle and the target is less than the distance d_2 between the vehicle and the threat. Under these circumstances, the *threat proximity* measure is equal to the distance between the target and the threat $d_R(V_i, G_j, T_2) = d_4$.

By contrast, if the distance between the vehicle and the threat is smaller than the distance between the vehicle and the target, then we estimate threat exposure using the minimum distance between the straight line path from vehicle to goal and the threat. This scenario is diagrammed in Figure 2 where the distance d_1 between the vehicle and T_1, is less than the distance d_B between the vehicle and the target. Consider the point on the straight line path from V_i to G_j that corresponds to the minimum distance from threat T_k; in Figure 2, this point occurs at the intersection of the dashed line and the straight line between vehicle and target. Then, the threat proximity is given by $d_R(V_i, G_j, T_1) = d_3$.

The *threat exposure* for pursuing target G_j is given by the minimum threat proximity over all threats

$$d_R(V_i, G_j) = \min_{T_k} d_R(V_i, G_j, T_k).$$

The rejectability of an assignment is high when threat exposure is high. Thus, following the same procedure as we did for the *ShortPath* heuristic, we can identify the rejectability function associated with the *AvoidThreats* heuristic,

$$\mu_{R_i}(G_j) = \frac{d_R(V_i, G_j) - \min_{G_j}\Big(d_R(V_i, G_j)\Big)}{\max_{G_j}\Big(d_R(V_i, G_j)\Big) - \min_{G_j}\Big(d_R(V_i, G_j)\Big)}. \tag{1}$$

Again, this function can be computed for all agents and targets in order MN time (times the number of threats), and assigns a rejectability of one to the target that exposes the vehicle to the most threat, and a rejectability of zero to the target that exposes the vehicle to the least threat.

4.3 Individually Satis£cing Assignments

For each vehicle, we can identify the set of targets that are individually satis£cing. This set is given by

$$S_{V_i} = \{G_j : \mu_{A_i}(G_j) \geq \mu_{R_i}(G_j)\}. \tag{2}$$

A vehicle balances the desire to take short paths with the desire to avoid threats; the path to a target must be short enough to justify the risks that arise by coming into proximity with threats. Since each vehicle must accept an assignment, these heuristics are encoded relative to the best and worst cases thereby causing actions to be evaluated with respect to the existing situation rather than with respect to an arti£cial standard.

5 Social Welfare

We have addressed how the two individual-based heuristics can be used to identify sets of individually satis£cing actions. We will combine the remaining two heuristics into a social welfare function to be maximized. Before we address how to construct this social welfare function, we explore the complexity of making group choices that are individually satis£cing.

5.1 Complexity Analysis

We are now in a position to explore the time complexity of searching through the set of all individually satis£cing options. Let $|S_{V_i}|$ denote the cardinality of the set S_{V_i}. The time complexity to £nd the optimal social welfare function over the set of satis£cing options is given by $\prod_{i=1}^{N} |S_{V_i}|$ which, in general, is much less than M^N since $|S_{V_i}| < M$.

To estimate how much complexity reduction is achieved, we ran 2000 trials with randomly selected positions for targets, threats, and vehicles. In an imaginary 1km by 1km world, targets and threats were randomly assigned locations within the world. The vehicles were assigned locations with uniform probability to a 200m border around the world. From the 2000 trials, the median complexity was 2400, the mean was 4437, the standard deviation of these complexities was 7214, and the maximum complexity was 122472. This indicates that, by far, the vast majority of circumstances produce complexities that are several orders of magnitude less than the optimal assignment problem, but there is a skew that causes some complexities to be only moderately better. However, only 9.3% of circumstances produce complexities greater than $10,000$ and, as discussed in the section on future work, most of these large search spaces occur under conditions where more heuristics could be applied to further narrow the search.

5.2 Constructing the Social Welfare

There are two competing objectives that affect how the team behavior is evaluated: the *MaxForce* heuristic drives the team toward assigning every vehicle to the same target, but the *MaxSpread* heuristic drives the team toward attacking as many targets as possible. We will address these objectives in turn.

From the *MaxForce* perspective, the more vehicles assigned to a target the better. We can represent this objective by assigning a value to the team size for each assigned target. In other words, we can identify a number that encodes how much better a larger team is than a smaller team. We use a monotonically increasing function that increases dramatically between one and three team members, indicating that a minimally acceptable team consists of two members. Let $\mathcal{G} = \{G_{j_1}, G_{j_2}, \ldots, G_{j_N}\}$ denote a set of targets assigned to each vehicle; G_{j_i} is the target assigned to vehicle V_i. The monotonic function that encodes the value of *MaxForce* is the sigmoid

$$U_{MaxForce}(G_j) = \frac{1}{1 + e^{-3(m(G_j)-2)}},$$

where $m(G_j) = |\{V_i : G_j = G_{j_i}\}|$ is the number of vehicles assigned to target G_j. We can use the value of the team size, encoded in $U_{MaxForce}(G_j)$, to estimate the value of an assignment. Let $\Gamma(\mathcal{G}) = \{G_j : \exists G_{j_i} \in \mathcal{G} \text{ for which } G_j = G_{j_i}\}$; $\Gamma(\mathcal{G})$ is the set of all targets that have vehicles assigned to them in the assignment $\mathcal{G}$. The value of this assignment from the perspective of the *MaxForce* heuristic is given by $V_{MaxForce}(\mathcal{G}) = \prod_{G_j \in \Gamma(\mathcal{G})} U_{MaxForce}(G_j)$. This value takes the product of each team-size value over the set of assigned target (the set of targets that have at least one vehicle attacking it). Using the product lets targets with small team sizes dramatically reduce the value of the overall assignment.

Turning attention to the *MaxSpread* heuristic, we note that the number of targets in the set $\Gamma(\mathcal{G})$ represents the spread of the assignment whence $V_{MaxSpread}(\mathcal{G}) = |\Gamma(\mathcal{G})|$.

Taking the product of $V_{MaxForce}$ and $V_{MaxSpread}$ yields the value of the assignment

$$V'(\mathcal{G}) = |\Gamma(\mathcal{G})| \times \prod_{G_j \in \Gamma(\mathcal{G})} U_{MaxForce}(G_j). \qquad (3)$$

Multiplying by the number of targets assigned allows the preference for large team sizes (i.e., *MaxForce*) to be tempered by the preference for having multiple targets attacked (i.e., *MaxSpread*). Since it is unlikely that all agents can feasibly attack the same target (it is unlikely that such an assignment is individually satis£cing), the product prevents undesirable assignments (with, for example, £ve agents assigned to one target and only one to another target) from receiving high preference. Thus, multiple, nearly equal-sized teams are preferred to multiple, differently-sized teams. $V'(\mathcal{G})$ for various team assignments is diagrammed in Figure 3.

One £nal addition signi£cantly improves the subjective quality of the resulting assignments. This improvement comes by letting the agent with the most dif£cult relative assignment act as a tie-breaker when choosing which target a team of a given size pursues. This is done by updating the value of the assignment as follows

$$V(\mathcal{G}) = V'(\Gamma(\mathcal{G})) + \epsilon \min_i [\mu_{A_i}(G_{j_i}) - \mu_{R_i}(G_{j_i})].$$

The multiplicative factor ϵ guarantees that team-size dominates the utility of an assignment, but that the vehicle that experiences the longest path or most exposure has some in¤uence over which target is assigned to its team. The assignment applied to the world is $\mathcal{G}^* = \arg\max V\mathcal{G})$.

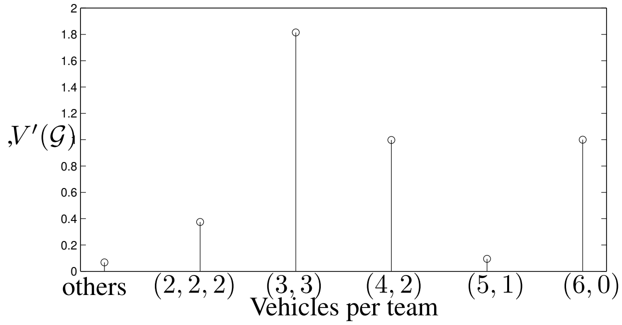

Figure 3: The value of different team assignments. The numbers in parenthesis indicates the number of vehicles assigned to each target; e.g., $(5,1)$ indicates that £ve vehicles are assigned to one target and one vehicle to another, and $(2,2,2)$ indicates that three targets are assigned to three, two-vehicle teams. Any team con£guration other than the ones shown (e.g. $(4,1,1)$, have values less than the value of the $(5,1)$ assignment.

6 Results

6.1 Qualitative Results

Figures 1 and 4-6 present snapshots of illustrative target assignments. The caption of each £gure discusses the unique characteristics of the assignment.

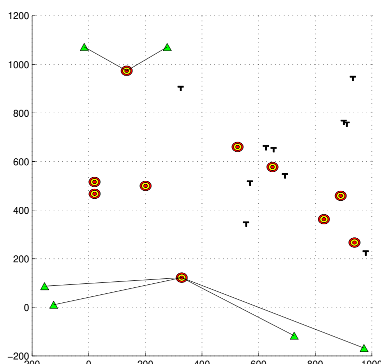

Figure 4: Notice how several vehicles converge on a relatively exposed target.

6.2 Average Team Size

One of the more interesting ways to quantify the algorithm behavior is in the average team size that is chosen. Figure 7 presents the results for 1000 exper-

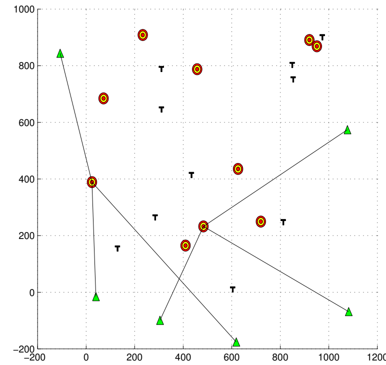

Figure 5: Notice how trajectories cross so that threats are avoided.

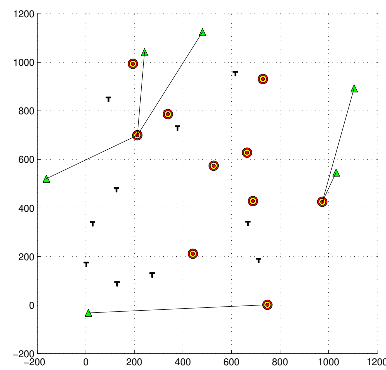

Figure 6: Notice how the vehicle on the Southwest corner has no threat-free path to rendezvous with the other vehicles. Consequently, it chooses the closest target.

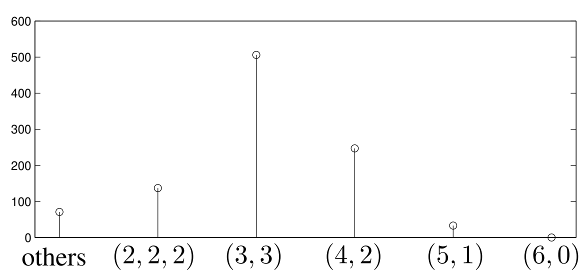

Figure 7: The frequency of occurrence of each team size for 1000 experimental trials.

imental trials. This £gure demonstrates that the algorithm tends to select teams with between two and four agents per team, and tends to avoid teams of one and teams greater than four. Thus, we can conclude that our team selection algorithm accomplishes our objective of maximizing the number of vehicles assigned to each target but also maximizing the spread of our attack. If vehicles are too far away to effectively coordinate (as enforced by the path length and threat exposure tradeoff), then they naturally form smaller strike forces to cover closer and more exposed targets.

7 Future Work

The heuristics that were used in this paper encode information that is necessary for making decisions. Different encodings should be explored to improve the usefulness of these heuristics. Of particular importance is improving the heuristic path of the vehicles because straightline trajectories are unreasonable. Future work should include using Voronoi-based paths as in [8], and then basing threat exposure and path-length on these paths.

Although the use of individually satis£cing choices constrains the size of the search space, the search is still exponential in complexity. We note that scenarios that yield large numbers of satis£cing assignments correspond to scenarios for which almost all assignments have equal evaluations . This implies that most assignments are valid whence additional heuristics can be used to further reduce the complexity of the search. Preliminary results using the non-domination principle from multi-attribute utility show promise. Additionally, iterative algorithms may provide a means for achieving polynomial complexity.

References

[1] R. W. Beard, W. C. Stirling, and R. L. Frost. A hierarchical coordination scheme for satellite formation initialization. In *AIAA Guidance, Navigation and Control Conference*, pages 677–685, Boston, MA, USA, August 1998.

[2] R. Borndörfer, A. Eisenblätter, M. Grötschel, and A. Martin. Frequency assignment in cellular phone networks. In *International Symposium on Mathematical Programming*, Lausanne, EPFL, August 24-29 1997.

[3] F. Brandt, W. Brauer, and G. Weiß. Task assignment in multiagent systems based on vickrey-type auctioning and leveled commitment contracting. In *CIA-2000 (Cooperative Information Agents) workshop proceedings*, Boston, USA, 2000.

[4] R. E. Burkard. Selected topics in assignment problems. Technical report, Institute of Mathematics B, Technical University Graz, Movember 1999. To appear in *Discrete Applied Mathematics*.

[5] M. A. Goodrich, W. C. Stirling, and E. R. Boer. Satis£cing revisited. *Minds and Machines*, 10:79–110, 2000.

[6] M. A. Goodrich, W. C. Stirling, and R. L. Frost. Model predictive satis£cing fuzzy logic control. *IEEE Transactions on Fuzzy Systems*, 7(3), June 1999.

[7] Yong Li, P. M. Pardalos, and M. G. C. Resende. A greedy randomized adaptive search procedure for the quadratic assignment problem. *DIMACS Series on Discrete Mathematics and Theoretical Computer Science*, 16:237–261, 1994.

[8] T. McLain and R. W. Beard. Cooperative rendezvous of multiple unmanned air vehicles. In *AIAA Guidance, Navigation and Control Conference*, Denver, CO, August 2000.

[9] H. A. Simon. *The Sciences of the Arti£cial*. MIT Press, 3rd edition, 1996.

[10] W. C. Stirling and M. A. Goodrich. Satis£cing games. *Information Sciences*, 114:255–280, March 1999.

[11] P. Stone and M. Veloso. Task decomposition and dynamic role assignment for real-time strategic teamwork. In *Intelligent Agents V - Proceedings of the Fifth International Workshop on Agent Theories, Architectures, and Languages (ATAL'98)*, Heidelberg, 1999. Springer Verlag.