



A Formal Analysis and Taxonomy of Task Allocation in Multi-Robot Systems

Brian Gerkey and Maja Matarić
Presentation by: Elaine Short




- Consider the problem of multi-robot task allocation (MRTA)
- Previous work has focused on creating working systems, rather than theoretical analysis.




- One idea common to all MRTA problems is the idea of utility, quality, and cost.
- We are trying to optimize by maximizing quality and minimizing cost.

$$U_{RT} = \begin{cases} Q_{RT} - C_{RT} & \text{if } R \text{ is capable of executing} \\ & T \text{ and } Q_{RT} > C_{RT} \\ 0 & \text{otherwise.} \end{cases}$$




- We will be using the idea of combinatorial optimization as a framework for analyzing these task assignment problems.
- Subset system:
 - A subset system (E, F) is a finite set of objects E and a non-empty collection F of subsets, called independent sets, of E that satisfies the property that if $X \in F$ and $\gamma \subseteq X$ then $\gamma \in F$.
- Subset maximization:
 - Given a subset system (E, F) and a utility function $u : E \rightarrow \mathbb{R}_+$, find an $X \in F$ that maximizes the total utility:

$$u(X) = \sum_{e \in X} u(e).$$




- Reorder the elements of $E = \{e_1, e_2, \dots, e_n\}$ such that: $u(e_1) \geq u(e_2) \geq \dots \geq u(e_n)$.
- Set $X := \emptyset$.
- For $j = 1$ to n : if $X \cup \{e_j\} \in F$ then $X = X \cup \{e_j\}$.

Note: The greedy algorithm will yield an optimal solution iff (E, F) is a matroid.




- Single Task (ST) vs. Multi-Task (MT)
- Single-robot Tasks (SR) vs. Multi-robot Tasks (MR)
- Instantaneous Assignment (IA) vs. Time-extended Assignment (TA)




- This is just an instance of the optimal assignment problem (OAP).
- Centralized algorithm (Hungarian method): $O(mn^2)$
- Distributed algorithm (Auction algorithm): “usually requiring time proportional to the maximum utility and inversely proportional to the minimum bidding increment”
- Central algorithm runs faster, but requires many more messages ($O(n^2)$ messages) than the distributed version (as few as $O(n)$).
- Practical note: Hungarian method is fast enough to be used in a control loop (even with old, slow computers).




Iterated Assignment




- This is a variant on the ST-SR-IA problem; imagine we have information coming in and need to reassign robots to tasks.
- Use Broadcast Local Eligibility (BLE):
 - If any robot remains unassigned, find the robot-task pair (i, j) with the highest utility. Otherwise, quit.
 - Assign robot i to task j and remove them from consideration.
 - Repeat until all robots have been assigned to tasks.
- This is a greedy algorithm, but OAP is not a matroid; this is 2-competitive.
- Practical note: useful for robot soccer, among other things.
- Practical note 2: the greedy algorithm may actually perform much better than 2-competitive for real-world problems.




- Assume robots that have already been reassigned cannot be reassigned (otherwise this is just iterated ST-SR-IA).
- MURDOCH Algorithm
 - When a new task is introduced, assign it to the most fit robot that is currently available.
- This algorithm is 3-competitive, and furthermore this is the best any on-line assignment algorithm can do (without further information).




- Build a time-extended model of robot assignments to minimize total cost.
- This is an instance of $R \parallel \sum w_j C_j$.
- Known to be strongly NP-hard.
- Approximation Algorithm:
 - Optimally solve the initial $m \times n$ assignment problem.
 - Use the Greedy algorithm to assign the remaining tasks in an on-line fashion, as the robots become available.
- This algorithm is 3-competitive (from previous results), but may perform much better in practice.
- There are other algorithms which accomplish this, but cannot be analyzed.




- Equivalent to dividing robots into task-specific groups, which is analogous to the Set Partitioning Problem (SPP).
- SPP is known to be strongly NP-hard, but is well studied and has good approximation algorithms.
- Open questions:
 - Are these approximations applicable to MRTA problems?
 - How to analyze the approximations?




- Need to consider all schedules for all coalitions.
- An instance of the multiprocessor scheduling problem: $MPT_m \parallel \sum w_j C_j$.
- This problem is strongly NP-hard, even with 2 processors.
- We could ignore the time-extended component, and solve greedily, but behavior of this algorithm is hard to analyze.



-
- Current robots are actuator-poor, thus this class of problems is uncommon.
 - Directly analogous to ST-MR-IA and ST-MR-TA, with tasks and robots switched.
 - Can be analyzed with the same tools.



- Can be recast as an instance of set cover, which is strongly NP-hard.
- Set cover can be approximated:
 - Chvátal (1979): logarithmic in the size of the largest feasible subset; running time is polynomial in the number of feasible subsets
 - Bar-Yehuda and Even (1981): maximum number of subsets to which any element belongs
- Practical note: this probably means that these algorithms would be best employed in situations where the potential coalitions are limited.
- Open question: how well do set cover algorithms work for MT-MR-A



- Contains an instance of the strongly-NP-hard problem $MPTmMPMn \parallel \sum w_j C_j$.
- No known approximations.




Table 1. Summary of Selected Iterated Assignment Architectures for MRTA

Name	Computation/ Iteration	Communication/ Iteration	Solution Quality
ALLIANCE ^a (Parker 1998)	$O(mn)$	$O(m)$	at least 2-competitive
BLE (Werger and Matarić 2001)	$O(mn)$	$O(mn)$	2-competitive
M+ (Botelho and Alami 1999)	$O(mn)$	$O(mn)$	2-competitive


^a In addition to solving the ST-SR-IA problem, the ALLIANCE architecture is also capable of building time-extended task schedules in order to solve a form of the ST-SR-TA problem (see Section 5.2.1).

Note. Shown here for each architecture are the computational and communication requirements, as well as solution quality.


Table 2. Summary of Selected Online Assignment Architectures for MRTA

Name	Computation/ Task	Communication/ Task	Solution Quality
MURDOCH (Gerkey and Matarić 2002b)	$O(1)$ / bidder $O(n)$ / auctioneer	$O(n)$	3-competitive
First-price auctions (Dias and Stentz 2001)	$O(1)$ / bidder $O(n)$ / auctioneer	$O(n)$	at least 3-competitive
Dynamic role assignment (Chaimowicz, Campos, and Kumar 2002)	$O(1)$ / bidder $O(n)$ / auctioneer	$O(n)$	at least 3-competitive


Note. Shown here for each architecture are the computational and communication requirements, as well as solution quality.



- Note the similarities: this suggests that the fundamental problem is similar, despite differences in the technical details.
- Analysis can also explain why these auction algorithms work well (solving optimizations).




-
- Other problems:
 - Interrelated utilities
 - Task constraints
 - Take advantage of domain-specific knowledge
 - Model how different types of “utility landscapes” affect the quality of the greedy solution.




Sequential Incremental- Value Auctions


Xiaoming Zheng and Sven Koenig




- We have a number of robots, and a number of targets, and want to visit all targets using our robots.
- Previous algorithms use an auction method (Sequential Single-Item auctions), this work will build on this idea, but will get better bounds on performance.



- Intuition:
 - Auctioneer releases the set of unassigned targets to the robots.
 - Each robot calculates the cost for it to get to each target (given the previous targets it was assigned), submits that as a bid.
 - The lowest cost-bid results in the assignment of one task to one robot.




```
1 function SSI-Auction ( $\bar{T}$ ,  $A$ )
2 inputs:  $\bar{T}$ : the set of targets  $T$ 
3            $A$ : the set of robots  $A$ 
4 outputs:  $\{T_a\}_{a \in A}$ : a complete assignment
5 for each robot  $a \in A$  do
6    $T_a \leftarrow \emptyset$ ;
7 while ( $\bar{T} \neq \emptyset$ ) do
8   /* Annunciation Stage */
9   the auctioneer announces  $\bar{T}$  to each robot  $a \in A$ ;
10  /* Bidding Stage */
11  for each robot  $a \in A$  do
12    for each target  $t \in \bar{T}$  do
13       $b_a^t \leftarrow c_a^{robot}(T_a \cup \{t\}) - c_a^{robot}(T_a)$ ;
14      robot  $a$  submits  $b_a^t$  to the auctioneer;
15    /* Winner-Determination Stage */
16     $(a, t) \leftarrow \arg \min_{(a \in A, t \in \bar{T})} b_a^t$ ;
17     $T_a \leftarrow T_a \cup \{t\}$ ;
18     $\bar{T} \leftarrow \bar{T} \setminus \{t\}$ ;
```



- Intuition:
 - Add another loop of rounds with an increasing bound b .
 - Auctioneer releases the list of unassigned targets, bound.
 - Eligible robots bid (eligible = robot which hasn't 'won' a round for this value of b).
 - Bid the number of targets the robot can reach in the bound b .
 - The robot with the winning bid (most targets) gets assigned those targets, and is removed from eligibility.




Sequential Incremental-Value Auction




```

1 function SIV-Auction ( $\bar{T}$ ,  $A$ ,  $b$ )
2 inputs:  $\bar{T}$ : the set of targets  $T$ 
3            $A$ : the set of robots  $A$ 
4            $b$ : a constant in  $(1, 2)$ 
5 outputs:  $\{T_a\}_{a \in A}$ : a complete assignment
6  $j \leftarrow 0$ ;
7 for each robot  $a \in A$  do
8    $T_a \leftarrow \emptyset$ ;
9 while ( $\bar{T} \neq \emptyset$ ) do
10   $j \leftarrow j + 1$ ;
11   $B \leftarrow b^{j+1}$ ;
12   $A' \leftarrow A$ ;
13  while ( $A' \neq \emptyset$ ) do
14    /* Annunciation Stage */
15    the auctioneer announces  $\bar{T}$  and  $B$  to each robot  $a \in A'$ ;
16    /* Bidding Stage */
17    for each robot  $a \in A'$  do
18       $T'_a \leftarrow \arg \max_{T' \subseteq \bar{T}: c_a^{path}(T') \leq B} |T'|$ ;
19      robot  $a$  submits  $T'_a$  to the auctioneer;
20    /* Winner-Determination Stage */
21     $a \leftarrow \arg \max_{a \in A'} |T'_a|$ ;
22     $A' \leftarrow A' \setminus \{a\}$ ;
23     $T_a \leftarrow T_a \cup T'_a$ ;
24     $\bar{T} \leftarrow \bar{T} \setminus T'_a$ ;

```

- Robots need to approximate the number of targets they can reach within the bound (this is NP-hard).
- Assume that we are doing this using a $(1/\alpha)$ -approximation.
- Practical note: these algorithms do exist.



Theorem 2 For all $j \geq 1$, $0.5 |T_j| \leq n_j - n_{j+1}$.

Corollary 1 For all $j \geq 1$, $n_{j+1} \leq 0.5 (n_j^* + n_j)$.


- $\{T_a^*\}_{a \in A}$: any complete assignment of the multi-robot routing problem and the order in which each robot should visit the targets assigned to it so that the resulting team cost is minimal (short: the optimal assignment) - if there is more than one, choose one arbitrarily;
- $c^* = \sum_{a \in A} c_a^{robot}(T_a^*)$: the team cost of the optimal assignment (short: minimal team cost);
- n_j^* : the number of targets whose latencies are larger than $0.5 \alpha b^{j+1}$ in the optimal assignment $\{T_a^*\}_{a \in A}$;
- n_j : the number of unassigned targets in the beginning of the j th round of an SIV auction; and
- T_j : the set of unassigned targets in the beginning of the j th round of the SIV auction whose latencies are at most $0.5 \alpha b^{j+1}$ in the optimal assignment.

Theorem 3 The team costs of SIV auctions are at most a factor of $O(1/\alpha)$ larger than minimal if each robot calculates its bids with a $(1/\alpha)$ -approximation algorithm for determining rooted k -MSTs.

Corollary 2 The team costs of SIV auctions are at most a factor of $O(1)$ larger than minimal if each robot calculates its bids with a constant factor approximation algorithm for determining rooted k -MSTs.



However...



- Constant-factor approximations for rooted k -MSTs are slow.
- Nearest neighbor algorithm will run fast, and is a good approximation.

Theorem 4 *The nearest-neighbor algorithm produces rooted k -trees whose costs are at most a factor of k larger than minimal.*

Corollary 3 *The team costs of SIV auctions are at most a factor of $O(|T|)$ larger than minimal if each robot calculates its bids with the nearest-neighbor algorithm for determining rooted k -MSTs (since $k \leq |T|$).*