16-311-Q Introduction to Robotics Fall'17

Lecture 29:
Multi-Robot Systems 2

Instructor:
Gianni A. Di Caro

جامعة كارنيجي ميلون في قطر
Carnegie Mellon University Qatar

ST-SR-IA: ONLINE ASSIGNMENT

- Tasks are revealed one at-a-time
- If robots can be *reassigned*, then solving each time the linear assignment provides the optimal solution, otherwise:

MURDOCH (2002)

  - When a new task is introduced, assign it to the most fit robot that is currently available.

- Greedy
- 3-competitive
- Performance bound is the best possible for any on-line assignment algorithm (Kalyana-sundaram, Pruhs 1993): without a model of the tasks that are to be introduced, and without the option of reassigning robots that have already been assigned, it is impossible to construct a better task allocator than MURDOCH.

ST-SR-TA: GENERALIZED ASSIGNMENT

$$
\begin{aligned}
\max \quad & \sum_{r=1}^{|R|} \sum_{t=1}^{|T|} U_{rt} x_{rt} \\
s.t. \quad & \sum_{t=1}^{|T|} c_{rt} x_{rt} \leq T_r \qquad r = 1, \ldots |R| \\
& \sum_{r=1}^{|R|} x_{rt} = 1 \qquad t = 1, \ldots |T| \\
& x_{rt} \in \{0, 1\}
\end{aligned}
$$

Robots get a schedule of tasks

More tasks than robots and the whole set should be assigned at the same time.
Future utilities are known

The “budget” constraints restricts the max number T_r of tasks (or the total time/energy to execute them based on some cost parameter c) that can be assigned to robot r

NP-hard!

ST-SR-TA: GENERALIZED ASSIGNMENT

$$
\begin{aligned}
\max \quad & \sum_{r=1}^{|R|} \sum_{t=1}^{|T|} U_{rt} x_{rt} \\
s.t. \quad & \sum_{t=1}^{|T|} c_{rt} x_{rt} \leq T_r \qquad r = 1, \ldots |R| \\
& \sum_{r=1}^{|R|} x_{rt} = 1 \qquad t = 1, \ldots |T| \\
& x_{rt} \in \{0, 1\}
\end{aligned}
$$

Approximated solution (not all tasks are jointly assigned):

1. Optimally solve the initial $R \times R$ assignment problem
2. Use the Greedy algorithm to assign the remaining tasks in an online fashion, as the robots become available.

Bound by 3-competitive greedy: as (|T|-|R|) goes to zero, gets optimal

ST-SR-TA: GENERALIZED ASSIGNMENT

If dependencies / constraints are included, “more” NP-Hard
→ If the utility is related to traveling distances the problem falls in the class of *m*TSP, VRP problems

Multi-robot routing

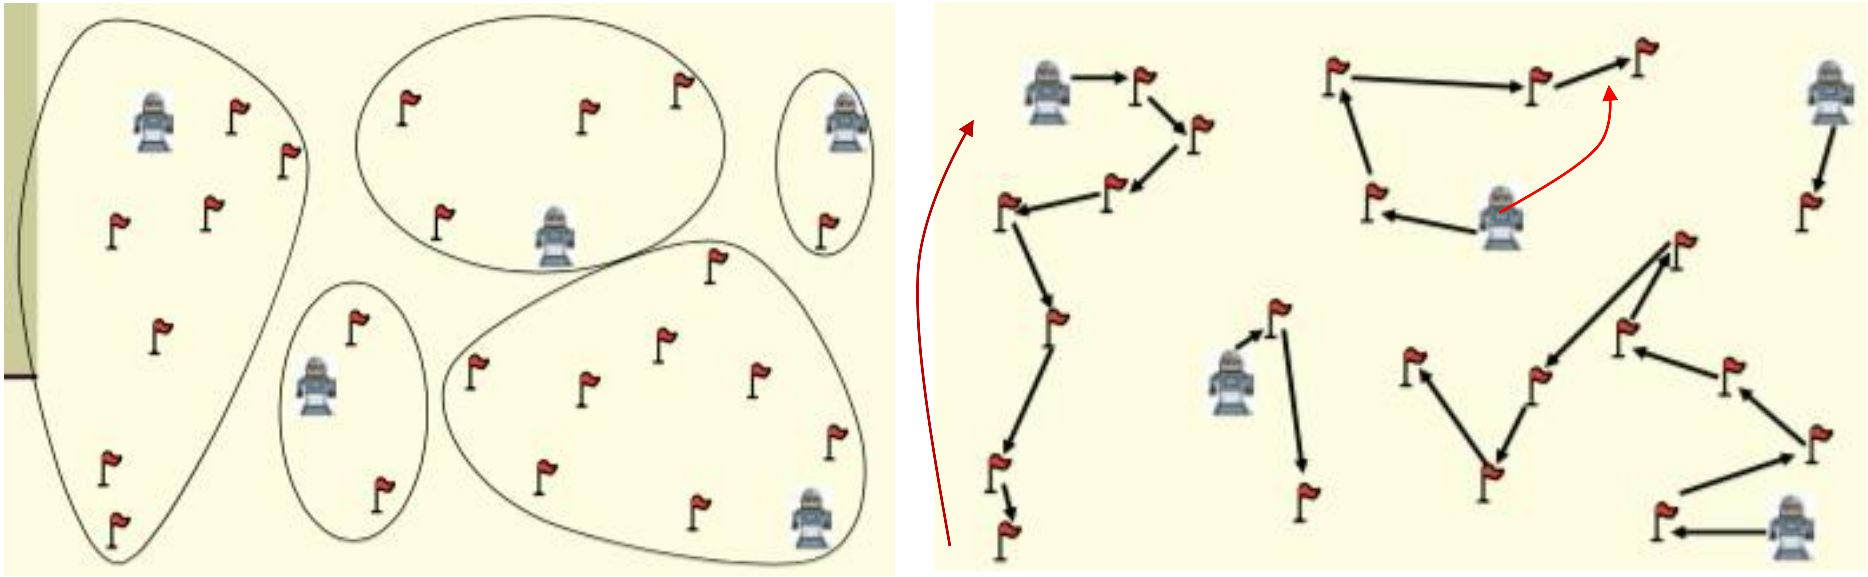

MT-SR-IA: GENERALIZED ASSIGNMENT

$$
\begin{aligned}
\max \quad & \sum_{r=1}^{|R|} \sum_{t=1}^{|T|} U_{rt} x_{rt} \\
s.t. \quad & \sum_{t=1}^{|T|} c_{rt} x_{rt} \leq T_r \qquad r = 1, \ldots |R| \\
& \sum_{r=1}^{|R|} x_{rt} = 1 \qquad t = 1, \ldots |T| \\
& x_{rt} \in \{0, 1\}
\end{aligned}
$$

Robots can work in ||
on multiple tasks

- The “capacity” constraint explicitly restricts the max number T_r of tasks that robot r can take, this time simultaneously
- Not common in the literature instances from MRTA

NP-hard!

MT-SR-TA: VRP

Robots can work in || on multiple tasks and
have a time-extended schedule of tasks
(quite uncommon in current MR literature)

Vehicle routing problems with capacity constraints and pick-up and delivery fall in this category:

- Multiple vehicles transporting multiple items (goods, people,…) and picking up items along the way
- Between a pick-up and delivery location the vehicle is dealing with MT
- Visiting multiple locations is equivalent to TA

NP-hard!

ST-MR-IA: SET PARTITIONING - COALITION FORMATION

- Model of the problem of dividing (partitioning) the set of robots into non-overlapping sub-teams (coalitions) to perform the given tasks *instantaneously assigned*
- This problem is mathematically equivalent to **set partitioning problem** in combinatorial optimization.

Cover (Partition) the elements in R (Robots) using the elements in CT (feasible coalition-task pairs) without duplicates (overlapping), and at the min cost / max utility

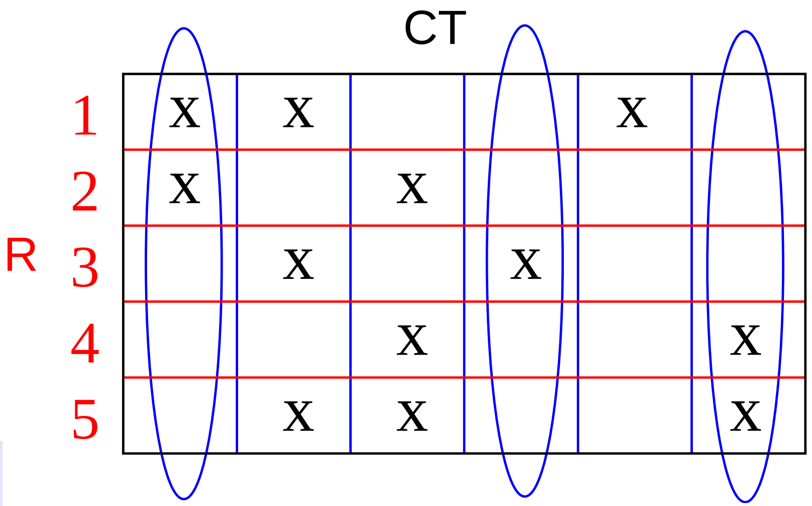

1	X	X			X	
2	X		X			
3		X		X		
4			X			X
5		X	X			X

$$\max / \min \; Z = \sum_{j=1}^{k} d_j x_j$$

$$s.t. \quad \sum_{j=1}^{k} a_{ij} x_j = 1, \quad \forall i = 1, \ldots, m$$

$$x_j \in \{0, 1\}, \quad \forall j = 1, \ldots, k$$

General SP model

NP-hard!

MT-MR-IA: SET COVERING - COALITION FORMATION

- Model of the problem of dividing (partitioning) the set of robots into sub-teams (coalitions) to perform the given tasks instantaneously assigned. Overlap is admitted to model MT, a robot can be in multiple coalitions
- This problem is mathematically equivalent to **set covering problem** in combinatorial optimization.

Cover (Partition) the elements in R (Robots) using the elements in CT (feasible coalition-task pairs) admitting duplicates (overlapping) and at the min cost / max utility

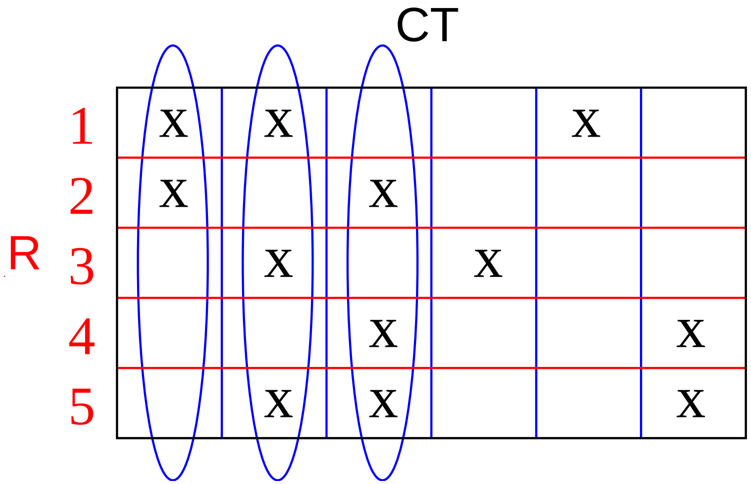

1	X	X			X	
2	X		X			
3		X		X		
4			X			X
5		X	X			X

$$\min\ Z = \sum_{j=1}^{k} c_j x_j$$

$$s.t. \quad \sum_{j=1}^{k} a_{ij} x_j \geqslant 1, \quad \forall i = 1, \dots, m$$

$$x_j \in \{0,1\}, \quad \forall j = 1, \dots, k$$

General SC model

NP-hard!

OTHER CASES

- ST-MR-TA: Involves both coalition formation and scheduling, and it’s mathematically equivalent to MT-SR-TA
- MT-MR-TA: Scheduling problem with multiprocessor tasks and multipurpose machines
- **Modeling of dependencies?** → G. Ayorkor Korsah, Anthony Stentz, and M. Bernardine Dias. 2013. A comprehensive taxonomy for multi-robot task allocation. Int. J. Rob. Res. 32, 12 (October 2013), 1495-1512.

SOLUTION APPROACHES

- Use the reference optimization models in a centralized scheme, solving the problems to optimality (e.g., Hungarian algorithm, IP solvers using branch-and-bound, optimization heuristics)
- Use the reference optimization models adopting a **top-down decentralized scheme** (e.g., all robots employ the same optimization model, and rely on local information exchange to build the model)
- Adopt different solution models *avoiding to explicitly formulate optimization problems.*
- Market-based approaches are an effective and popular option
- Emergent/Swarm approaches: effective / simpler alternative

MARKET-BASED: BASIC IDEAS

- Based on the economic model of a free market
- Each robot seeks to maximize individual “profit”
- Individual profit helps the common good
- An auctioneer (i.e. a robot spotting a new task) offers tasks (or roles, or resources) in an announcement phase
- Robots can negotiate and bid for tasks based on their (estimated) utility function
- Once all bids are received or the deadline has passed, the auction is cleared in the winner determination phase: the auctioneer decides which items to award and to whom.
- Decisions are made locally but effects approach optimality
  - Preserve advantages of distributed approach

MARKET-BASED: BASIC IDEAS

- Robots model an economy:
  - Accomplish task → Receive revenue
  - Consume resources → Incur cost
  - Robot goal: maximize own profit
  - Trade tasks and resources over the market (auctions)
- By maximizing individual profits, team finds better solution
- Time permitting → more centralized
- Limited computational resources → more distributed

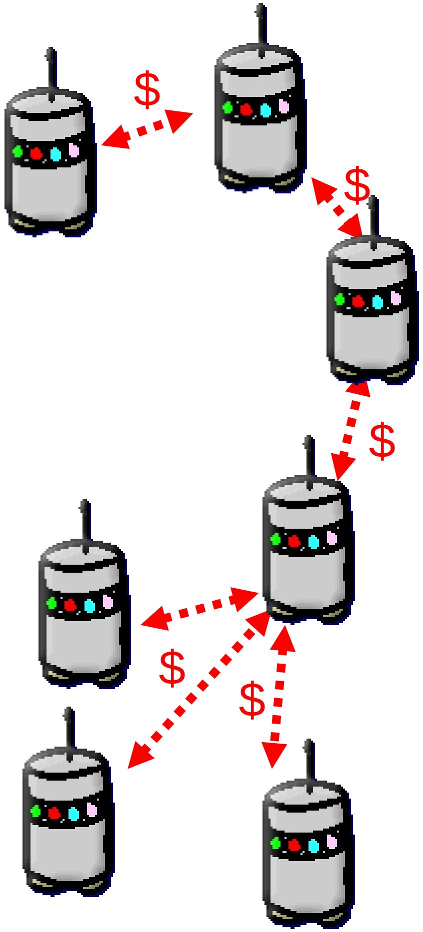

MARKET-BASED: BASIC IDEAS

- Utility = $Revenue\ - Cost$
- Team revenue is sum of individual revenues
- Team cost is sum of individual costs
- Costs and revenues set up per application
    - Maximizing individual profits must move team towards globally optimal solution
- Robots that produce well at low cost receive a larger share of the overall profit

MARKET-BASED: IMPLEMENTATIONS

- MURDOCH (Gerkey and Matarić, IEEE Trans. On Robotics and Automation, 2002 / IJRR 2004)
- M+ (Botelho and Alami, ICRA 1999)
- TraderBots (Dias et al., multiple publications 1999-2006)

Basic Ideas of Emergent TA

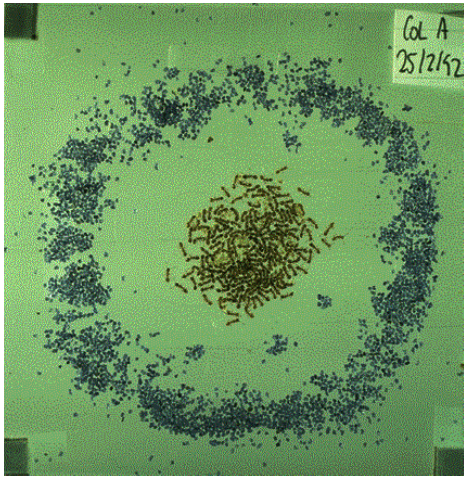

Ideas and models from clustering and labor division behaviors in ant colonies

Brood care:

- Larvae are sorted in such a way that different brood stage are arranged in concentric rings
- Smaller larvae are in the center, larger larvae on the periphery

Cemetery organization:

- Clustering corpses to form cemeteries
- Each ants seems to move randomly while picking up or depositing (dropping) corpses
- Pick up or drop: decision based on local information
- The combination of these very simple behaviors from individual ants give raise to the emergence of colony-level complex behaviors of cluster formation

TASK ALLOCATION BASED ON RESPONSE THRESHOLD

- Response thresholds refer to the likelihood of reacting to task-associated stimuli (e.g. the presence of a corps or a larva, the height of a pile of dirty dishes to wash)
- Individuals with a low threshold perform a task at a lower level of stimulus than individuals with high thresholds
- Individuals become engaged in a specific task when the level of task-associated stimuli exceeds their thresholds
- If a task is not performed by individuals, the intensity of the corresponding stimulus increases
- Intensity decreases as more ants (agents) perform the task
- The task-associated stimuli serve as stigmergic variable

SINGLE TASK ALLOCATION

- Let s_j be the intensity of task-j-associated stimuli
- A response threshold, θ_{kj}, determines the tendency of individual k to respond to the stimulus, s_j, associated with task j
- Individual k engages in task j with probability

$$P_{\theta_{kj}}(s_j) = \frac{s_j^{\omega}}{s_j^{\omega} + \theta_{kj}^{\omega}}$$

where $\omega > 1$ determines the steepness of the threshold

- For $s_j << \theta_{kj}$, $P_{\theta_{kj}}(s_j)$ is close to zero, and the probability of performing task j is very small
- For $s_j >> \theta_{kj}$, the probability of performing task j is close to one

SINGLE TASK ALLOCATION

- Assume only one task
- The probability that an inactive ant will become active

$$P(\vartheta_k = 0 \rightarrow \vartheta_k = 1) = \frac{s^2}{s^2 + \theta_k^2}$$

  $\vartheta_k = 0$ indicates that ant k is inactive

  $\vartheta_k = 1$ indicates that the ant is performing the task

- An active ant spends an average $1/p$ time performing the task
- Change in stimulus intensity:

$$s(t+1) = s(t) + \sigma - \gamma n_{act}$$

  σ is the increase in demand

  γ is the decrease associated with one ant performing the task

  n_{act} is the number of active ants

SINGLE TO MULTIPLE TASK ALLOCATION

- The more ants engaged in the task, the smaller the intensity, s, and consequently, the smaller the probability that an inactive ant will take up the task
- If there are not enough ants busy with the task (i.e. $\sigma > \gamma n_{act}$), the probability increases that inactive ants will participate in the task
- Multiple tasks:
    - Let there be n_j tasks
    - Let n_{kj} be the number of workers of caste k performing task j
    - Each individual has a vector of thresholds, θ_k
    - After $1/p$ time units of performing task j, the ant stops with this task, and selects another task

SUMMARY

- Characteristics and basic taxonomy of multi-agents systems
- Taxonomy of multi-robot task allocation (MRTA) problems
- Optimization models for the different classes of MRTA problems
- Computational complexity of the different classes
- Basic solution approaches exploiting the optimization models
- Basic ideas about market-based methods
- Basic ideas about ant-based task allocation