Dynamic Vehicle Routing for Robotic Networks
Lecture #8: Different Task Models

Francesco Bullo[1] Emilio Frazzoli[2] Marco Pavone[2]
Ketan Savla[2] Stephen L. Smith[2]

[1]CCDC
University of California, Santa Barbara
bullo@engineering.ucsb.edu

[2]LIDS and CSAIL
Massachusetts Institute of Technology
{frazzoli,pavone,ksavla,slsmith}@mit.edu

Workshop at the 2010 American Control Conference
Baltimore, Maryland, USA, June 29, 2010, 8:30am to 5:00pm

Motivation for Team Forming

- Group of vehicles monitoring a region
- Several different **sensing modalities**:
  - electro-optical,
  - infra-red,
  - synthetic aperture radar,
  - foliage penetrating radar,
  - etc.
- Each event requires a **subset of sensing modalities**
- Equip each vehicle with a single sensing modality
- **Form appropriate team** to properly assess each event

How do we create teams in real-time to observe each event (service each request)?

Lecture outline

Literature Review

Scaling laws in Robotic Networks

V. Sharma, M. Savchenko, E. Frazzoli, and P. Voulgaris. Transfer time complexity of conflict-free vehicle routing with no communications. *International Journal of Robotics Research*, 26(3):255–272, 2007

S. L. Smith and F. Bullo. Monotonic target assignment for robotic networks. *IEEE Transactions on Automatic Control*, 54(9):2042–2057, 2009

Throughput vs Delay in Wireless Networks

G. Sharma, R. Mazumdar, and N. Shroff. Delay and capacity trade-offs in mobile ad hoc networks: A global perspective. In *IEEE Conf. on Computer Communications*, pages 1–12, Barcelona, Spain, April 2006

A. El Gamal, J. Mammen, B. Prabhakar, and D. Shah. Optimal throughput-delay scaling in wireless networks. Part I: The fluid model. *IEEE Transactions on Information Theory*, 52(6):2568–2592, 2006

Graph Coloring

T. A. McKee and F. R. McMorris. *Topics in Intersection Graph Theory*, volume 2 of *Monographs on Discrete Mathematics and Applications*. SIAM, 1999

B. Korte and J. Vygen. *Combinatorial Optimization: Theory and Algorithms*, volume 21 of *Algorithmics and Combinatorics*. Springer, 4 edition, 2007

Lecture outline

1. Dynamic Team Forming
2. Three Policies
3. Analysis of Policies

Dynamic Team Forming

Set of services $\{r_1, \ldots, r_k\}$.

Vehicle properties:

- k different vehicle types.
- Vehicle type $j \in \{1, \ldots, k\}$, can provide only service r_j.

Task (demand) model:

- Poisson and Uniform arrivals
- Each task requires a subset of services in $\{r_1, \ldots, r_k\}$.
- $\mathcal{K}$ different types of tasks
- Tasks of type α arrive at rate λ_α
- Task completed when required vehicles **simultaneously** spend on-site service time at location.

S. L. Smith and F. Bullo. The dynamic team forming problem: Throughput and delay for unbiased policies. *Systems & Control Letters*, 58(10-11):709–715, 2009

Load Factor and Stability

- $R_\alpha \in \{0,1\}^k$ is zero-one column vector recording services required for task-type α.
- on-site service for task-type α is $\bar{s}_\alpha$
- m_j vehicles provide service r_j.

Necessary stability condition:

$$[R_1 \ \cdots \ R_{\mathcal{K}}] \begin{bmatrix} \lambda_1 \bar{s}_1 \\ \vdots \\ \lambda_{\mathcal{K}} \bar{s}_{\mathcal{K}} \end{bmatrix} < \begin{bmatrix} m_1 \\ \vdots \\ m_k \end{bmatrix}$$

Load factor is now a vector

Example of Team Forming

- $k = 4$ different services, $\{r_1, r_2, r_3, r_4\}$.
- $m = 8$ vehicles, two of each type: $m_j = 2$ for $j \in \{1, 2, 3, 4\}$.
- $\mathcal{K} = 6$ task types, $\{r_1\}, \{r_2\}, \{r_3\}, \{r_4\}, \{r_1, r_3\}, \{r_2, r_4\}$.

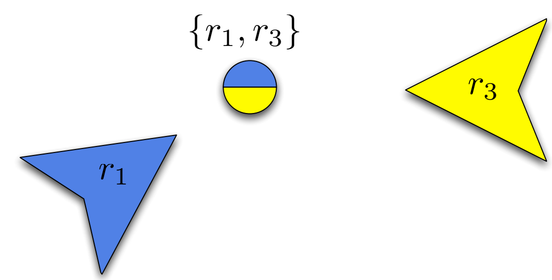

Task type $\alpha = \{r_1, r_3\}$ has on-site service $\bar{s}_\alpha$, arrival rate λ_α, and

$$R_\alpha = [1\ 0\ 1\ 0]^\top.$$

Lecture outline

1. Dynamic Team Forming
2. Three Policies
   - Complete Team
   - Task-Specific Team Policy
   - Scheduled Task-Specific Team Policy
3. Analysis of Policies

Task-Type Unbiased Policies

For a policy π:

- System time of each task-type $\overline{T}_{\pi,1}, \dots, \overline{T}_{\pi,\mathcal{K}}$
- Feasible set of system times are subset of $\mathbb{R}^{\mathcal{K}}$
- Optimization space similar to priority queues, but with teaming

To simplify, consider **task-type unbiased policies**

$$\overline{T}_{\pi,1} = \overline{T}_{\pi,2} = \cdots = \overline{T}_{\pi,\mathcal{K}} =: \overline{T}_{\pi}$$

and the optimization: $\inf_{\pi} \overline{T}_{\pi}$.

Policy 1: Complete Team Policy

Complete Team Policy

1: Form $\min\{m_1, \dots m_k\}$ teams of k vehicles, each team contains one vehicle of each type.
2: Have each team meet and move as a single entity.
3: In each region run UTSP policy (from Lecture 3).

Can also use Divide & Conquer policy for each team

Policy 1: Complete Team Policy

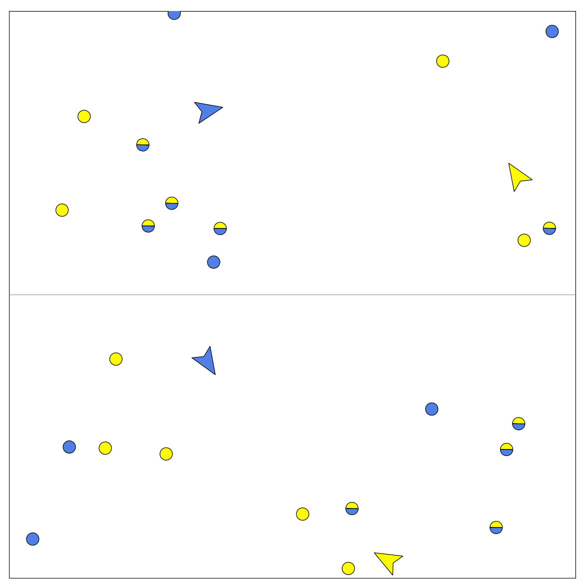

- Two services y, b
- 3 task-types $y, b, \{y, b\}$.
- 4 vehicles
  - 2 yellow
  - 2 blue

Policy 2: Task-specific Team Policy

- m_j vehicles provide service r_j.
- r_j appears in $e_j^\top [R_1 \cdots R_\mathcal{K}] \mathbf{1}_\mathcal{K}$ task types.
- If $m_j \geq e_j^\top [R_1 \cdots R_\mathcal{K}] \mathbf{1}_\mathcal{K} \Rightarrow$ enough vehicles of type j to create dedicated team for each task type.
- Create m_{TST} teams, where:

$$m_{\text{TST}} := \left\lfloor \min_j \left\{ \frac{m_j}{e_j^T R \mathbf{1}_\mathcal{K}} \right\} \right\rfloor$$

Task-Specific Team Policy

1: For each of the $\mathcal{K}$ task types, create m_{TST} teams of vehicles.
2: Service each task by one of its m_{TST} corresponding teams, according to the UTSP policy.

Policy 2: Task-Specific Team Policy

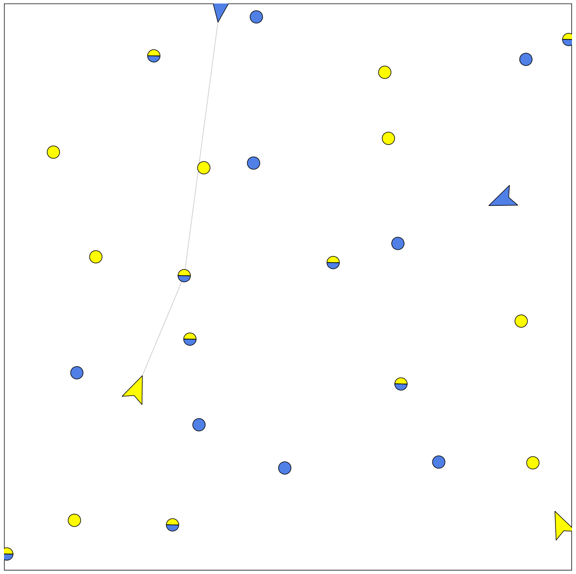

- task types: $\{y\}, \{b\}, \{y, b\}$
- two vehicles of each type
- y, b each appear in two task-types

$$m_{\text{TST}} = 1$$

Policy 3: Preliminary Result

Definition (Service schedule)

A partition of task types into L time slots, such that:

- each type appears in exactly one time slot, and
- task types in each time slot are pairwise disjoint.

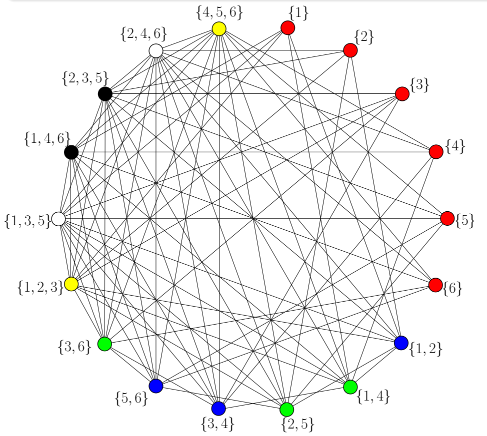

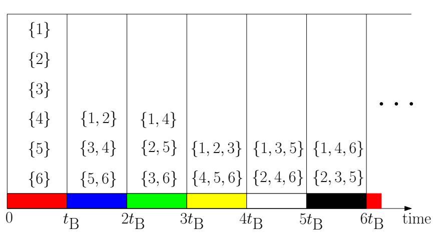

Policy 3: Scheduled Task-Specific Team Policy

Scheduled Task-specific team policy

Partition $\mathcal{Q}$ into $\min_i\{m_i\}$ regions and assign one robot of each type to each region.

1: In each region form a queue for each task type.
2: For each time slot in the schedule:
   1. Divide robots into teams to form required task types.
   2. For each team, service corresponding queue with TSP tour.

Policy 3: Scheduled Task-Specific Team Policy

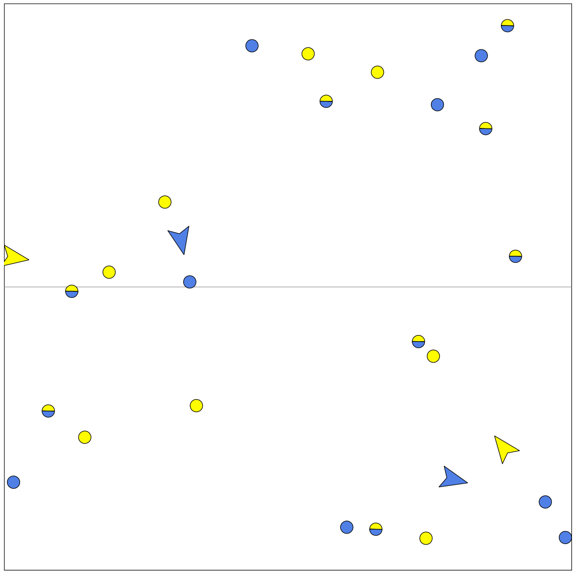

Service schedule:

- two time slots $L = 2$
- slot one: $\{y\}, \{b\}$
- slot two: $\{y, b\}$

Lecture outline

Assumptions for Analysis

Assumptions:

1. $m_i = m/k$ for each vehicle type i.
2. $\lambda_\alpha = \lambda/\mathcal{K}$ for each task-type α.
3. on-site service has mean $\bar{s}$ and is upper bounded by $s_{\max}$.
4. $p\mathcal{K}$ of the $\mathcal{K}$ task-types require service r_j, where $p \in [1/k, 1]$.

- With assumptions, necessary stability condition becomes

$$\frac{\lambda}{m} < \frac{1}{pk\bar{s}}.$$

- Define **per-vehicle throughput** as $B_m := \lambda/m$.

Throughput vs System Time Profile

$$B_m \mapsto \begin{cases} \max\left\{\overline{T}_{\min}, \dfrac{\overline{T}_{\mathrm{ord}}(B_m/B_{\mathrm{crit}})}{(1 - B_m/B_{\mathrm{crit}})^2}\right\}, & \text{if } B_m < B_{\mathrm{crit}}, \\ +\infty, & \text{if } B_m \geq B_{\mathrm{crit}}. \end{cases}$$

- $\overline{T}_{\min}$ = minimum achievable system time for positive throughput.
- B_{crit} = maximum achievable throughput (or capacity).
- $\overline{T}_{\mathrm{ord}}$ = system time at a constant fraction of capacity $(3 - \sqrt{5})/2 \approx 38\%$ of capacity B_{crit}.

Example (Single vehicle DVR)

$B_{\mathrm{crit}} = 1/\bar{s}$

$\overline{T}_{\min} = \mathbb{E}_\varphi[\|X - p^*\|]/v + \bar{s}$ (light load)

$\overline{T}_{\mathrm{ord}} \approx C(\int_Q \varphi^{1/2}(x)dx)^2/v^2$ (heavy load numerator)

Throughput vs System Time Profile

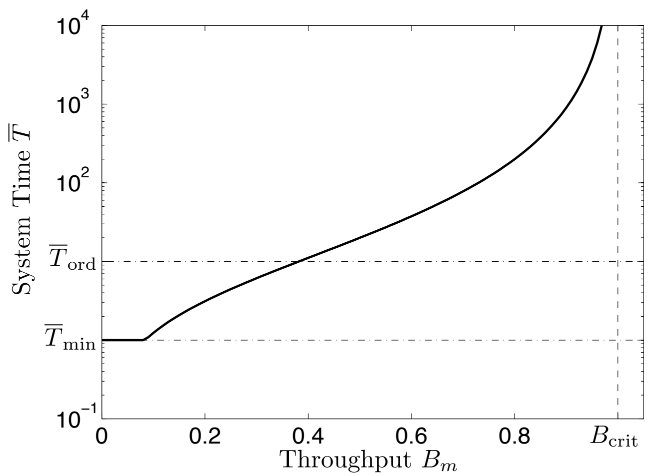

System Time for each Policy

	$\overline{T}_{\rm min}$	$\overline{T}_{\rm ord}$	$B_{\rm crit}$
Lower bound ($\overline{T}^*$)	$\sqrt{k}$	k	$\frac{1}{pk\bar{s}}$
Complete Team	$\sqrt{k}$	k	$\frac{1}{k\bar{s}}$
Task-Specific	$\sqrt{pk\mathcal{K}}$	$pk\mathcal{K}$	$\frac{1}{pk\bar{s}}$
Scheduled Task-Specific	$L\sqrt{k}$	Lk	$\frac{\mathcal{K}}{s_{\max}Lk}$

where $L \in [p\mathcal{K}, \mathcal{K}]$

Best policies for different scenarios:

- If throughput is low, then use complete team
- If p is close to 1, then use complete team
- If p is close to $1/k$, then for best capacity use
  - Task-Specific if enough vehicles
  - Scheduled Task-Specific otherwise

Lecture outline

1. Dynamic Team Forming
2. Three Policies
   - Complete Team
   - Task-Specific Team Policy
   - Scheduled Task-Specific Team Policy
3. Analysis of Policies
   - Throughput vs System Time
   - Comparison of Results

Workshop Structure and Schedule

8:00-8:30am	*Coffee Break*	
8:30-9:00am	Lecture #1:	Intro to dynamic vehicle routing
9:05-9:50am	Lecture #2:	Prelims: graphs, TSPs and queues
9:55-10:40am	Lecture #3:	The single-vehicle DVR problem
10:40-11:00am	*Break*	
11:00-11:45pm	Lecture #4:	The multi-vehicle DVR problem
11:45-1:10pm	*Lunch Break*	
1:10-2:10pm	Lecture #5:	Extensions to vehicle networks
2:15-3:00pm	Lecture #6:	Extensions to different demand models
3:00-3:20pm	*Coffee Break*	
3:20-4:20pm	Lecture #7:	Extensions to different vehicle models
4:25-4:40pm	Lecture #8:	Extensions to different task models
4:45-5:00pm		Final open-floor discussion