On Partitioning Policies in Dynamic Vehicle Routing Problems

Emilio Frazzoli

Laboratory for Information and Decision Systems
Aeronautics and Astronautics Department
Massachusetts Institute of Technology

thanks to
A. Arsie, J. Enright, M. Pavone, K. Savla *and* **F. Bullo** *(UCSB)*

Swarms Workshop

Block Island, RI

June 4, 2009

ARES
Aerospace Robotics and Embedded Systems Laboratory

A class of dynamic vehicle routing problems

- Spatially-localized tasks are generated over time by an exogenous process in a geographic area of interest

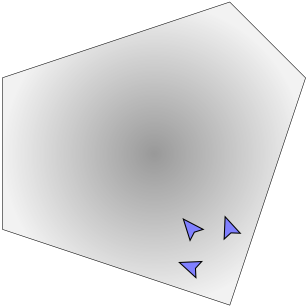

- Mobile agents can complete tasks by moving to the tasks' locations
- Optimize a Quality of Service measure
- Cooperation occurs via **workload sharing**

A class of dynamic vehicle routing problems

- Spatially-localized tasks are generated over time by an exogenous process in a geographic area of interest

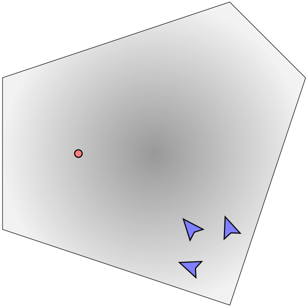

- Mobile agents can complete tasks by moving to the tasks' locations
- Optimize a Quality of Service measure
- Cooperation occurs via **workload sharing**

A class of dynamic vehicle routing problems

- Spatially-localized tasks are generated over time by an exogenous process in a geographic area of interest
- Mobile agents can complete tasks by moving to the tasks' locations
- Optimize a Quality of Service measure
- Cooperation occurs via **workload sharing**

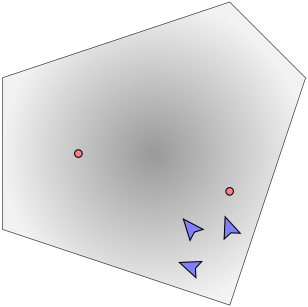

A class of dynamic vehicle routing problems

- Spatially-localized tasks are generated over time by an exogenous process in a geographic area of interest

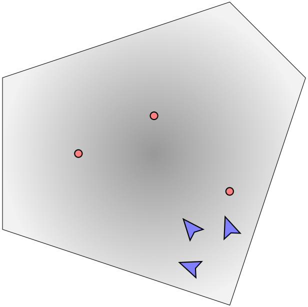

- Mobile agents can complete tasks by moving to the tasks' locations
- Optimize a Quality of Service measure
- Cooperation occurs via **workload sharing**

A class of dynamic vehicle routing problems

- Spatially-localized tasks are generated over time by an exogenous process in a geographic area of interest
- Mobile agents can complete tasks by moving to the tasks' locations
- Optimize a Quality of Service measure
- Cooperation occurs via **workload sharing**

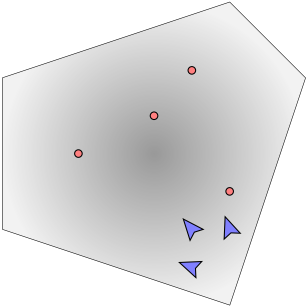

A class of dynamic vehicle routing problems

- Spatially-localized tasks are generated over time by an exogenous process in a geographic area of interest

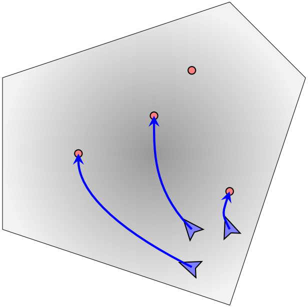

- Mobile agents can complete tasks by moving to the tasks' locations
- Optimize a Quality of Service measure
- Cooperation occurs via **workload sharing**

A class of dynamic vehicle routing problems

- Spatially-localized tasks are generated over time by an exogenous process in a geographic area of interest

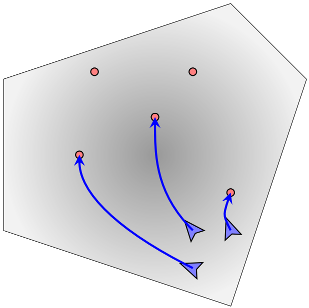

- Mobile agents can complete tasks by moving to the tasks' locations
- Optimize a Quality of Service measure
- Cooperation occurs via **workload sharing**

A class of dynamic vehicle routing problems

- Spatially-localized tasks are generated over time by an exogenous process in a geographic area of interest

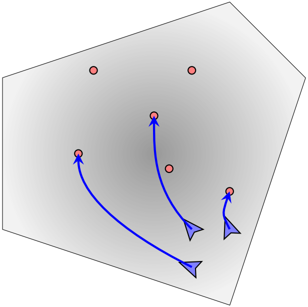

- Mobile agents can complete tasks by moving to the tasks' locations

- Optimize a Quality of Service measure

- Cooperation occurs via **workload sharing**

A class of dynamic vehicle routing problems

- Spatially-localized tasks are generated over time by an exogenous process in a geographic area of interest
  For example, spatio-temporal Poisson point process, time intensity λ, spatial distribution φ supported over $\mathcal{Q}$
- Mobile agents can complete tasks by moving to the tasks' locations

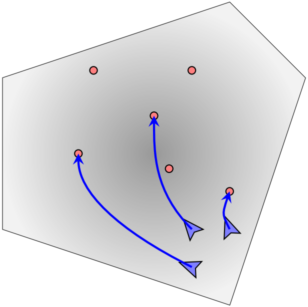

- Optimize a Quality of Service measure
- Cooperation occurs via **workload sharing**

A class of dynamic vehicle routing problems

- Spatially-localized tasks are generated over time by an exogenous process in a geographic area of interest
  For example, spatio-temporal Poisson point process, time intensity λ, spatial distribution φ supported over $\mathcal{Q}$
- Mobile agents can complete tasks by moving to the tasks' locations
  For simplicity, consider m identical agents with infinite capacity
- Optimize a Quality of Service measure

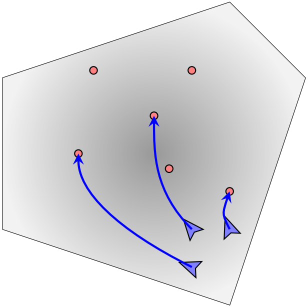

- Cooperation occurs via **workload sharing**

A class of dynamic vehicle routing problems

- Spatially-localized tasks are generated over time by an exogenous process in a geographic area of interest
  For example, spatio-temporal Poisson point process, time intensity λ, spatial distribution φ supported over $\mathcal{Q}$
- Mobile agents can complete tasks by moving to the tasks' locations
  For simplicity, consider m identical agents with infinite capacity
- Optimize a Quality of Service measure
  - *The delay between task generation and task completion*
  - *Percentage of tasks completed before expiration*
- Cooperation occurs via **workload sharing**

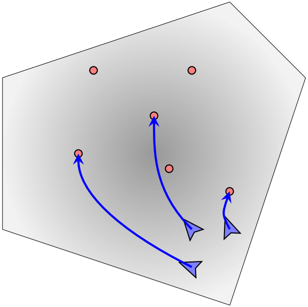

A class of dynamic vehicle routing problems

- Spatially-localized tasks are generated over time by an exogenous process in a geographic area of interest
  For example, spatio-temporal Poisson point process, time intensity λ, spatial distribution φ supported over $\mathcal{Q}$
- Mobile agents can complete tasks by moving to the tasks' locations
  For simplicity, consider m identical agents with infinite capacity
- Optimize a Quality of Service measure
  - *The delay between task generation and task completion*
  - *Percentage of tasks completed before expiration*
- Cooperation occurs via **workload sharing**

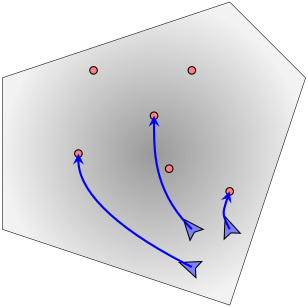

Dynamic Traveling Repairperson Problem
[Psaraftis '88, Bertsimas *et al.* '90s]

Dynamic Vehicle Routing: challenges

- Control policies:
  - Task assignment
  - Task scheduling
  - “Loitering” policies.

Dynamic Vehicle Routing: challenges

- Control policies:
  - Task assignment
  - Task scheduling
  - "Loitering" policies.
- Constraints:
  - Limited on-board computational resources
  - Limited/unreliable communications and sensing
  - Algebraic/differential/integral constraints on the vehicles' motion

Dynamic Vehicle Routing: challenges

- Control policies:
  - Task assignment
  - Task scheduling
  - "Loitering" policies.
- Constraints:
  - Limited on-board computational resources
  - Limited/unreliable communications and sensing
  - Algebraic/differential/integral constraints on the vehicles' motion
- Analysis objectives:
  - Provable performance guarantees.
  - Performance as a function of system parameters

The Euclidean, light load case

- Let us first consider the following basic case:
  - The target generation rate is very small: $\lambda \to 0^+$.
  - The target spatial distribution is supported on a convex, compact set $\mathcal{Q}$ (i.e., $\varphi(q) > 0 \Leftrightarrow q \in \mathcal{Q}$.)
  - Agents can move with bounded speed V.
- In such case:
  - With high probability all vehicles will have enough time to return to some "loitering" station between task completion/generation times.
  - The problem is reduced to the choice of the loitering stations $g^* = (g_1^*, g_2^*, \ldots, g_m^*) \in \mathcal{Q}^m$ that minimizes the system time.

The Weber (or multi-median) problem

- The optimal loitering station placement minimizes the continuous Weber, or multi-median function:

$$H_m(g) = \int_{\mathcal{Q}} \min_{i \in \{1,\ldots,m\}} \|g_i - q\|_2 \, \varphi(q) \; dq$$

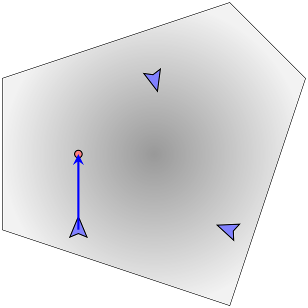

The Weber (or multi-median) problem

- The optimal loitering station placement minimizes the continuous Weber, or multi-median function:

$$H_m(g) = \int_{\mathcal{Q}} \min_{i \in \{1,\ldots,m\}} \|g_i - q\|_2 \, \varphi(q) \, dq$$
$$= \sum_{i=1}^{m} \int_{\mathcal{V}_i(g)} \|g_i - q\|_2 \, \varphi(q) \, dq.$$

- The optimal task assignment is based on the *Voronoi partition* generated by the loitering stations.

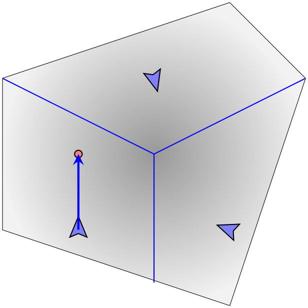

The Weber (or multi-median) problem

- The optimal loitering station placement minimizes the continuous Weber, or multi-median function:

$$H_m(g) = \int_{\mathcal{Q}} \min_{i \in \{1,\dots,m\}} \|g_i - q\|_2 \, \varphi(q) \, dq$$
$$= \sum_{i=1}^{m} \int_{\mathcal{V}_i(g)} \|g_i - q\|_2 \, \varphi(q) \, dq.$$

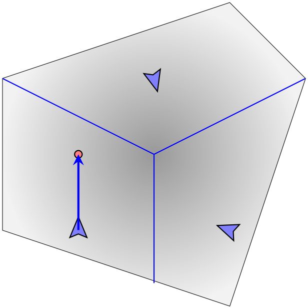

- The optimal task assignment is based on the *Voronoi partition* generated by the loitering stations.
- The continuous Weber function is differentiable (as long as the points g_i are distinct) but not convex $\Rightarrow$ NP-hardness.

The Weber (or multi-median) problem

- The optimal loitering station placement minimizes the continuous Weber, or multi-median function:

$$H_m(g) = \int_{\mathcal{Q}} \min_{i \in \{1,\ldots,m\}} \|g_i - q\|_2 \, \varphi(q) \, dq$$
$$= \sum_{i=1}^{m} \int_{\mathcal{V}_i(g)} \|g_i - q\|_2 \, \varphi(q) \, dq.$$

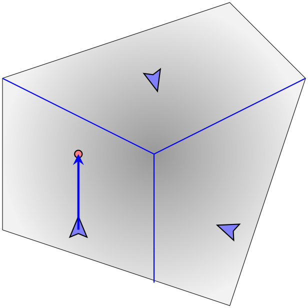

- The optimal task assignment is based on the *Voronoi partition* generated by the loitering stations.
- The continuous Weber function is differentiable (as long as the points g_i are distinct) but not convex $\Rightarrow$ NP-hardness.
- The optimal performance scales as $O(1/\sqrt{m})$.

Decentralized control laws

Both these control policies are known to converge to critical points of the performance function

Lloyd-like **[Cortes et al., '02]**

- Move their loitering station towards the median of their own region.
- Service targets in own region, returning to the loitering station when done.

- Needs (i) communication with neighbors, (ii) knowledge of spatial density φ.

MacQueen-like **[Arsie et al., '07]**

- If there is an outstanding task: Move towards the nearest task location.
- Otherwise: Move to a position that minimizes the average distance to the locations of tasks previously completed by the agent.

- No communication between agents, no need to know the spatial density φ, collaboration in learning as well as workload sharing.
- Slower convergence.

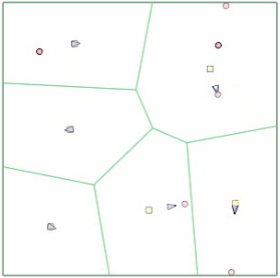

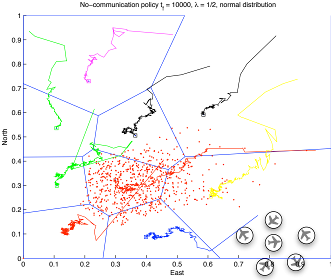

Introducing Differential Constraints

- What if the vehicle dynamics are more complex?
  E.g., airplanes, moving on paths with bounded curvature
- The problem can be restated:
  - The target generation rate is very small: $\lambda \to 0^+$.
  - The target spatial distribution is uniform on a convex, compact set $\mathcal{Q}$.
  - Vehicles move with fixed speed V, on paths with curvature bounded by $1/\rho$.
- In such case:
  - Vehicles cannot stop.
  - Strategies are more complex than defining a "loitering point."
- How many of the results from the Euclidean case carry over to this case?

The Median Circling (MC) policy

Control policy

Assign "virtual" generators to each agent. All agents do the following, in parallel (possibly asynchronously):

- Update the position of their generator according to a gradient descent law.
- Service targets in own region, returning to a "loitering circle" of radius 2.91ρ centered on their generator position when done.

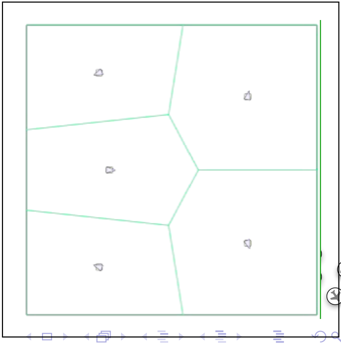

- We have

$$\lim_{\lambda \to 0^+} T_{\mathrm{MC}} \leq \frac{H_m^*(\mathcal{Q}) + 3.76\rho}{V} = T^* + \frac{c3.76\rho}{V}$$

- Furthermore,

$$\lim_{H_m^* \to +\infty, \lambda \to 0^+} \frac{T_{\mathrm{MC}}}{T^*} = 1,$$

$$\lim_{H_m^* \to 0^+, \lambda \to 0^+} T_{\mathrm{MC}} \geq \left(\frac{\pi}{2} - \frac{2}{\pi}\right) \frac{\rho}{V}$$

The Strip Loitering (SL) policy

Control policy

- Design a closed path P containing parallel segments crossing the environment $\mathcal{Q}$, at a distance w. All agents move along this path, equally spaced. When a new target arrives, the closest agent (taking the dynamics into account) is responsible for visiting it. Optimize over w.

- We have that

$$\lim_{H_m \to 0^+, \lambda \to 0^+} T_{\mathrm{SL}} m^{1/3} \leq \frac{1.238}{V} (\rho W H + 10.38 \rho^2 H)^{1/3}.$$

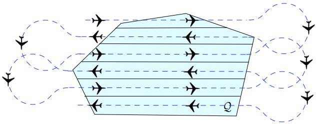

- Furthermore,

$$\lim_{H_m^* \to 0^+, \lambda \to 0^+} T_{\mathrm{SL}} = 0$$

$$\lim_{H_m^* \to 0^+, \lambda \to 0^+} \frac{T_{\mathrm{SL}}}{T^*} = \alpha,$$

Phase transition

- We have two policies: Median Circling (MC), and Strip Loitering (SL). Which is better?
- Define the non-holonomic density $d_\rho = \rho^2 m/A$.
  - MC is optimal when $d_\rho \to 0$, but any strategy based on separate regions of responsibility is arbitrarily bad when $d_\rho \to +\infty$.
  - Conversely for SL.
- The optimal organization changes from territorial (MC) to "gregarious" (SL) depending on the "non-holonomic density" of the agents.
- **Endogenous phase transition**, only depends on internal system characteristics, not on external stimuli (assuming uniform target density).

Estimate of the critical density

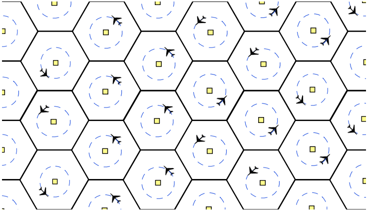

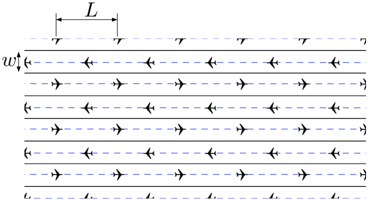

- Ignoring boundary conditions (e.g., consider the infinite plane), we can compute the system time for each policy exactly:

$$T_{\mathrm{SL}} < T_{\mathrm{MC}} \quad \Leftrightarrow \quad d_\rho > 0.0587$$

(i.e., transition occurs when the area of the dominance region is about 4-5 times the area of the minimum turning radius circle).

- Simulation results in a finite environment yielded a critical density $d_\rho = 0.0925$.
- (SL less attractive because of U-turns)

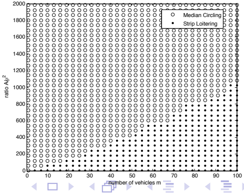

Conclusions and future work

- Dynamic Vehicle Routing: a broad and interesting class of problems combining (differential) geometry, combinatorial optimization, queueing theory.
  - adaptive and distributed policies
  - game-theoretic formulations
  - Demand models: priorities, impatience, dynamics.
  - Task models: pick-up and delivery, team formation, humans in the loop.
  - Sensing models: limited sensing range, dynamic data harvesting, environmental sensing.
- Differential constraints may introduce fundamental changes in the problem and in the solution: phase transitions may appear, dictating optimal strategies as a function of (collective) system's parameters.
  - Design mechanisms for individuals to detect whether a phase transition should occur.
  - Applications to biological systems?