Dynamic Vehicle Routing for Robotic Networks
Lecture #6: Different Demand 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: Time-Critical Tasks

Motivating Scenario

- Group of UAVs equipped with sensors, **monitoring region**
- Alerted of events that require **close-range observation**

Events with **time constraints:**

- Each event must be observed within a time-window

Events with **priority levels:**

- Each event has associated level of importance (e.g. 1 to 10)

Lecture outline

Lecture outline

M. Pavone and E. Frazzoli. Dynamic vehicle routing with stochastic time constraints. In *IEEE Int. Conf. on Robotics and Automation*, Anchorage, AK, May 2010

M. Pavone, N. Bisnik, E. Frazzoli, and V. Isler. A stochastic and dynamic vehicle routing problem with time windows and customer impatience. *ACM/Springer Journal of Mobile Networks and Applications*, 14(3):350–364, 2009

DVR with stochastic time constraints

Model:

- basic DVR model +
- demand j **active** for a random patience time G_j
- G_j's i.i.d. sequence $\sim F_G$
- demand j **expires** if not serviced within G_j

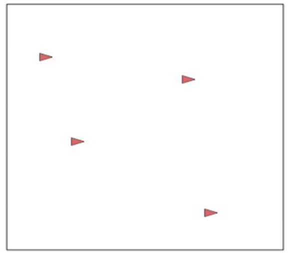

Service constraint:

- $\lim_{j\to+\infty} \mathbb{P}_\pi[W_j < G_j]$: acceptance probability for policy π
- $\phi^{\mathrm{d}} \in (0,1)$: desired acceptance probability
- **constraint**: $\lim_{j\to+\infty} \mathbb{P}_\pi[W_j < G_j] \geq \phi^{\mathrm{d}}$

Problem formulation

Problem statement

Solve problem OPT:

$$\min_{\pi} \quad |\pi|, \quad \text{subject to} \quad \lim_{j\to\infty} \mathbb{P}_\pi[W_j < G_j] \geq \phi^{\mathrm{d}}$$

Well-posedness

- Existence: $\lim_{j\to\infty} \mathbb{P}_\pi[W_j < G_j]$ exists for all π
- Ergodicity: $\lim_{j\to\infty} \mathbb{P}_\pi[W_j < G_j] = \lim_{t\to+\infty} N^s(t)/N(t)$ (a.s.)

Proof sketch:

- main idea: theory of regenerative processes
- regeneration points: times a new demand finds the system empty
- expected length of busy cycles is finite
- use classic limit theorems

Problem formulation

Problem statement

Solve problem OPT:

$$\min_{\pi} \quad |\pi|, \quad \text{subject to} \quad \lim_{j\to\infty} \mathbb{P}_\pi[W_j < G_j] \geq \phi^{\mathrm{d}}$$

Well-posedness

- Existence: $\lim_{j\to\infty} \mathbb{P}_\pi[W_j < G_j]$ exists for all π
- Ergodicity: $\lim_{j\to\infty} \mathbb{P}_\pi[W_j < G_j] = \lim_{t\to+\infty} N^s(t)/N(t)$ (a.s.)

Proof sketch:

- main idea: theory of regenerative processes
- regeneration points: times a new demand finds the system empty
- expected length of busy cycles is finite
- use classic limit theorems

Lower bound

Intuition for lower bound:

$$\mathbb{P}[W_j < G_j] \leq \mathbb{P}\left[\min_{k\in\{1,\dots,m\}} \frac{\|X_j - X_k\|}{v} < G_j\right] \leq \sup_{(p_1,\dots,p_m)\in\mathcal{Q}^m} \underbrace{\mathbb{P}\left[\min_{k\in\{1,\dots,m\}} \frac{\|X_j - p_k\|}{v} < G_j\right]}_{\doteq \mathcal{H}(p_1,\dots,p_m)}$$

Lower bound

OPT is lower bounded by:

$$\underline{OPT}: \min_{m\in\mathbb{N}_{>0}} \quad m \qquad \text{s.t.} \quad \sup_{(p_1,\dots,p_m)\in\mathcal{Q}^m} \mathcal{H}(p_1,\dots,p_m) \geq \phi^{\mathrm{d}}$$

Devised algorithms to solve $\underline{OPT}$

Lower bound

Intuition for lower bound:

$$\mathbb{P}[W_j < G_j] \leq \mathbb{P}\left[\min_{k\in\{1,\dots,m\}} \frac{\|X_j - X_k\|}{v} < G_j\right]$$
$$\leq \sup_{(p_1,\dots,p_m)\in\mathcal{Q}^m} \underbrace{\mathbb{P}\left[\min_{k\in\{1,\dots,m\}} \frac{\|X_j - p_k\|}{v} < G_j\right]}_{\doteq \mathcal{H}(p_1,\dots,p_m)}$$

Lower bound

OPT is lower bounded by:

$$\underline{OPT}: \min_{m\in\mathbb{N}_{>0}} \quad m$$
$$\text{s.t.} \quad \sup_{(p_1,\dots,p_m)\in\mathcal{Q}^m} \mathcal{H}(p_1,\dots,p_m) \geq \phi^{\mathrm{d}}$$

Devised algorithms to solve $\underline{OPT}$

Lower bound

Intuition for lower bound:

$$\mathbb{P}[W_j < G_j] \leq \mathbb{P}\left[\min_{k\in\{1,\dots,m\}} \frac{\|X_j - X_k\|}{v} < G_j\right]$$
$$\leq \sup_{(p_1,\dots,p_m)\in\mathcal{Q}^m} \underbrace{\mathbb{P}\left[\min_{k\in\{1,\dots,m\}} \frac{\|X_j - p_k\|}{v} < G_j\right]}_{\doteq \mathcal{H}(p_1,\dots,p_m)}$$

Lower bound

OPT is lower bounded by:

$$\underline{OPT}: \min_{m\in\mathbb{N}_{>0}} \quad m$$
$$\text{s.t.} \quad \sup_{(p_1,\dots,p_m)\in\mathcal{Q}^m} \mathcal{H}(p_1,\dots,p_m) \geq \phi^{\mathrm{d}}$$

Devised algorithms to solve $\underline{OPT}$

NDA policy (optimal as $\lambda \to 0$)

Nearest Depot Assignment (NDA) policy

```
Compute maximum of H:  (p̄_1, ..., p̄_m).
Then:
 1: p̄_k is depot of kth vehicle
 2: nearest-depot assignment
 3: FCFS service
```

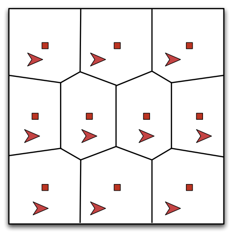

Proof sketch:

- as usual, as $\lambda \to 0^+$, the problem reduces to optimal pre-positioning

NDA policy (optimal as $\lambda \to 0$)

Nearest Depot Assignment (NDA) policy

```
Compute maximum of H:  (p̄_1, ..., p̄_m).
Then:
 1: p̄_k is depot of kth vehicle
 2: nearest-depot assignment
 3: FCFS service
```

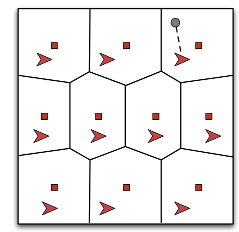

Proof sketch:

- as usual, as $\lambda \to 0^+$, the problem reduces to optimal pre-positioning

NDA policy (optimal as $\lambda \to 0$)

Nearest Depot Assignment (NDA) policy

Compute maximum of $\mathcal{H}$: $(\bar{p}_1, \dots, \bar{p}_m)$.
Then:

1: $\bar{p}_k$ is depot of kth vehicle
2: nearest-depot assignment
3: FCFS service

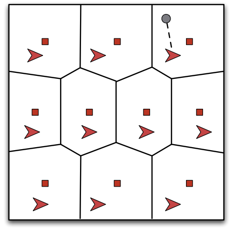

Proof sketch:

- as usual, as $\lambda \to 0^+$, the problem reduces to optimal pre-positioning

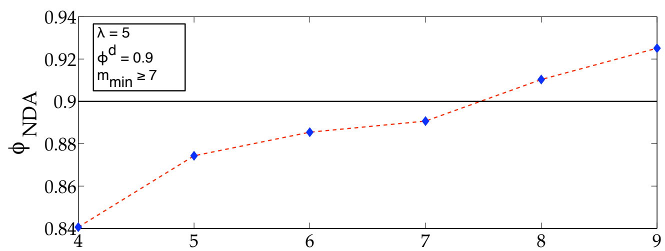

Batch policy

Batch (B) policy

Partition $\mathcal{Q}$ into m simultaneously equitable subregions and assign one vehicle to each subregion. Then:

1: each vehicle services demands by forming TSP tours

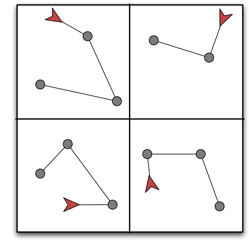

Performance of batch policy

- if s=0: $m_B = \min\left\{ m \,\middle|\, \sup_{\theta \in \mathbb{R}_{>0}} (1 - F_G(\theta))(1 - \frac{\lambda \cdot \mathrm{const}}{\theta m^2}) \geq \phi^{\mathrm{d}} \right\}$
- with time windows: $m_B / m^* \leq 3.78$, when λ large and $\phi^d \to 1^-$

Characterization of batch policy

Proof sketch (m=1):

- upper bound expected length of TSP tour with $\mathrm{const} \cdot \lambda / m^2$, via control-theoretical methods
- use Markov's ineq to lower bound:

$$\mathbb{P}[W < G] \geq \mathbb{P}[W < G | 2\,\mathrm{TSP} < \theta] \mathbb{P}[2\,\mathrm{TSP} < \theta]$$
$$\geq (1 - F_G(\theta))(1 - \mathbb{E}[2\,\mathrm{TSP}]/\theta)$$

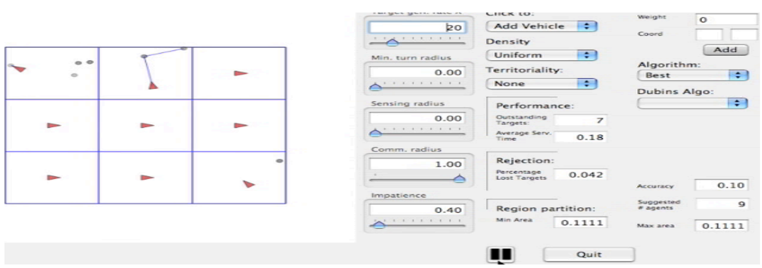

Lecture outline

Demands with priority levels

- m vehicles
- n classes of demands
  - $1 =$ highest priority
  - $n =$ lowest priority
- Poisson arrivals $\lambda_1, \dots, \lambda_n$
- locations uniformly distributed
  can extend to non-uniform φ

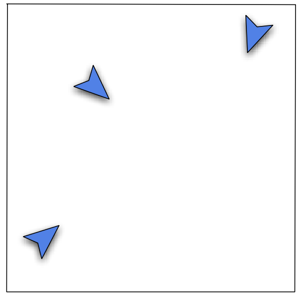

Steady-state expected system-time $\overline{T}_1, \dots, \overline{T}_n$

Goal for vehicles

Minimize $c_1\overline{T}_1 + \dots + c_n\overline{T}_n$ $\quad (\uparrow c_i \Rightarrow \uparrow$ priority of class $i)$

S. L. Smith, M. Pavone, F. Bullo, and E. Frazzoli. Dynamic vehicle routing with priority classes of stochastic demands. *SIAM Journal on Control and Optimization*, 48(5):3224–3245, 2010

Demands with priority levels

- m vehicles
- n classes of demands
  - $1 =$ highest priority
  - $n =$ lowest priority
- Poisson arrivals $\lambda_1, \dots, \lambda_n$
- locations uniformly distributed
  can extend to non-uniform φ

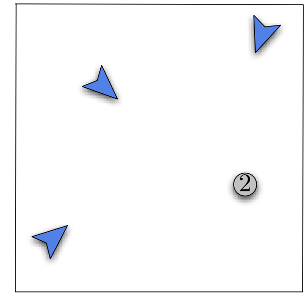

Steady-state expected system-time $\overline{T}_1, \dots, \overline{T}_n$

Goal for vehicles

Minimize $c_1\overline{T}_1 + \dots + c_n\overline{T}_n$ $\quad (\uparrow c_i \Rightarrow \uparrow$ priority of class $i)$

S. L. Smith, M. Pavone, F. Bullo, and E. Frazzoli. Dynamic vehicle routing with priority classes of stochastic demands. *SIAM Journal on Control and Optimization*, 48(5):3224–3245, 2010

Demands with priority levels

- m vehicles
- n classes of demands
  - $1 =$ highest priority
  - $n =$ lowest priority
- Poisson arrivals $\lambda_1, \dots, \lambda_n$
- locations uniformly distributed
  can extend to non-uniform φ

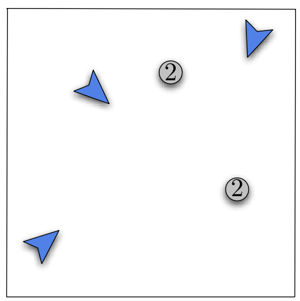

Steady-state expected system-time $\overline{T}_1, \dots, \overline{T}_n$

Goal for vehicles

Minimize $c_1\overline{T}_1 + \dots + c_n\overline{T}_n$ $\quad (\uparrow c_i \Rightarrow \uparrow$ priority of class $i)$

S. L. Smith, M. Pavone, F. Bullo, and E. Frazzoli. Dynamic vehicle routing with priority classes of stochastic demands. *SIAM Journal on Control and Optimization*, 48(5):3224–3245, 2010

Demands with priority levels

- m vehicles
- n classes of demands
  - $1 =$ highest priority
  - $n =$ lowest priority
- Poisson arrivals $\lambda_1, \dots, \lambda_n$
- locations uniformly distributed
  can extend to non-uniform φ

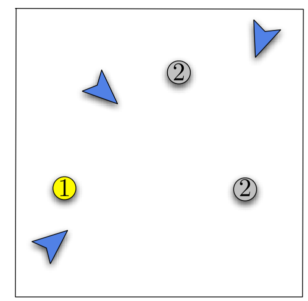

Steady-state expected system-time $\overline{T}_1, \dots, \overline{T}_n$

Goal for vehicles

Minimize $c_1\overline{T}_1 + \dots + c_n\overline{T}_n$ $\quad (\uparrow c_i \Rightarrow \uparrow$ priority of class $i)$

S. L. Smith, M. Pavone, F. Bullo, and E. Frazzoli. Dynamic vehicle routing with priority classes of stochastic demands. *SIAM Journal on Control and Optimization*, 48(5):3224–3245, 2010

Literature Review

Classic Priority Queueing

L. Kleinrock. *Queueing Systems. Volume II: Computer Applications*. Wiley, New York, 1976

E. G. Coffman Jr. and I. Mitrani. A characterization of waiting time performance realizable by single-server queues. *Operations Research*, 28(3):810–821, 1980

Related Combinatorial Problems

A. Blum, P. Chalasani, D. Coppersmith, B. Pulleyblank, P. Raghavan, and M. Sudan. The minimum latency problem. In *ACM Symposium on the Theory of Computing*, pages 163–171, Montreal, Canada, 1994

M. Z. Spivey and W. B. Powell. The dynamic assignment problem. *Transportation Science*, 38(4):399–419, 2004

A. Blum, S. Chawla, D. R. Karger, T. Lane, A. Meyerson, and M. Minkoff. Approximation algorithms for orienteering and discounted-reward TSP. *SIAM Journal on Computing*, 37(2):653–670, 2007

Load Factor and Stability

Stable: Queue remains bounded

Define **load factor** as

$$\varrho := \frac{\lambda_1 \bar{s}_1 + \cdots + \lambda_n \bar{s}_n}{m}$$

- λ_i = arrival rate for class i
- $\bar{s}_i$ = average on-site service time for class i

As before, **necessary stability condition** is $\varrho < 1$

Two asymptotic regimes

1. Light load $\varrho \to 0^+$
2. Heavy load $\varrho \to 1^-$

Light load

In light load:

- Each vehicle can return to a median between arrivals
- Priority levels do not change behavior.

Optimal solution:

m vehicle SQM policy is optimal (or an adaptive policy)

m Stochastic Queueing Median (m-SQM)

```
Compute m-median locations and assign one vehicle to
each location.
Then:
 1: service demands in FCFS order
 2: return to median after each service is completed
```


Lower Bound in Heavy Load

Let $\overline{T}_c^*$ = optimal value of cost $c_1\overline{T}_1 + \cdots + c_n\overline{T}_n$.

Lower bound for every policy

$$\overline{T}_c^* \geq \frac{\beta_{\mathrm{TSP}}|Q|}{2m^2v^2(1-\varrho)^2} \sum_{\alpha=1}^{n} \left(c_\alpha + 2 \sum_{j=\alpha+1}^{n} c_j \right) \lambda_\alpha$$

Problem parameters:

- arrival rates $\lambda_1, \ldots, \lambda_n$
- weights $c_1, \ldots, c_n$
- number of vehicles m
- environment area $|Q|$
- vehicle speed v

Proof Idea of Lower Bound

- Allow **remote service** of some classes: $r_\alpha \in \{0,1\}$ for each class α
- travel distance is $r_\alpha \bar{d}_\alpha$

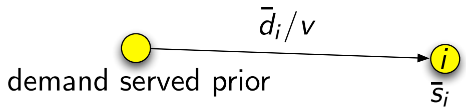

- For stability: $\sum_{i=1}^{n} \lambda_i \left(r_i \bar{d}_i/v + \bar{s}_i\right) < m$
- Can bound travel distance as

$$\bar{d}_\alpha \geq \frac{\beta_{\mathrm{TSP}}}{\sqrt{2}} \sqrt{\frac{|Q|}{\sum_i r_i \overline{N}_i}}$$

- generates a linear program with $2^n - 1$ constraints, one for each combination $\{r_1, \dots, r_n\}$
- solution to LP is largest lower bound

Separate Queues Policy

Input: Probability distribution $\mathbf{p} = [p_1, \dots, p_n]$.

Separate Queues Policy

```
Partition environment into m equal area regions and
assign one vehicle to each region.
Then:
 1: Select a class according to probability dist p
 2: Service all demands of selected class following TSP
 3: Repeat
```

Policy performance optimized over $\mathbf{p}$.

Separate Queues Performance

Heavy load performance

For the SQ policy,

$$\frac{\overline{T}_{c,\mathrm{SQ}}}{\overline{T}_c^*} \leq 2n^2$$

as $\varrho \to 1^-$.

- $n =$ number of classes
- independent of $\varrho, c, \bar{s}, \lambda$

Heuristic Improvements:

1. Receding horizon: service only a fraction η of TSP
2. when following TSP, service newly arrived demands within ϵ of TSP.

$$\epsilon \sqrt{\frac{\mu |Q|}{\sum_{\alpha=1}^{n} \overline{N}_\alpha}},$$

where μ is fractional in tour length (i.e., 0.1 for 10% increase)

Simulation of Separate Queues Policy

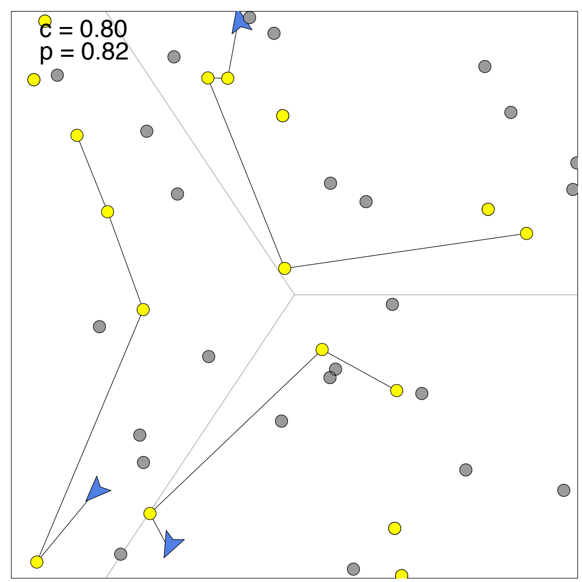

Simulation:

- class 1 = yellow
- class 2 = grey
- $c_1 = 0.8$ and $c_2 = 0.2$
- $\mathbf{p} = [0.82, 0.18]$

Proof idea for upper bound

- In heavy-load, shortest path through N points:

$$= \beta_{\mathrm{TSP}} \sqrt{|Q| N} \quad \text{with prob. 1 (BHH theorem)}$$

- Study expected # of outstanding demands at each iteration

$$N_i(t+1) \leq f\big(N_1(t), \ldots, N_m(t), \mathbf{p}, \lambda, \bar{s}\big)$$

- Function f has a linear part plus a sub-linear part
- Bound evolution by stable linear system for all $\varrho < 1$

$$\mathcal{N}(t+1) = A(\mathbf{p}, \lambda, \bar{s})\mathcal{N}(t) + B(\mathbf{p}, \lambda, \bar{s})$$

- Allows computation of $\limsup_{t \to +\infty} \mathcal{N}_i(t)$
- Apply Little's theorem $\overline{N}_i = \lambda_i \overline{T}_i$

Lecture outline

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