Dynamic Vehicle Routing for Robotic Networks
Lecture #3: The single-vehicle DVR problem

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

Lecture outline

D. J. Bertsimas and G. J. van Ryzin. A stochastic and dynamic vehicle routing problem in the Euclidean plane. *Operations Research*, 39:601–615, 1991

The problem

DVR - distinct features

- service demands vary over time
- information about future is stochastic
- real-time routing policies
- queueing phenomena

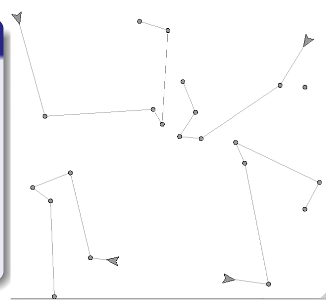

DVR is fundamentally **a queueing problem**:

1. arrival process
2. service model
3. performance measure

General queueing-theoretical model for DVR 1/2

Arrival process: spatio-temporal Poisson

1. time intensity $\lambda > 0$
2. spatial density φ: $\mathbb{P}\,[\text{demand in } \mathcal{S}] = \int_{\mathcal{S}} \varphi(x)\, dx$
3. inter-arrival times and locations are i.i.d.

Service model:

1. m holonomic vehicles with maximum velocity v
2. vehicles provide a random on-site service
3. on-site service times are i.i.d. (equal on average to $\bar{s}$)
4. demand removed from the system upon on-site service completion

General queueing-theoretical model for DVR 1/2

Arrival process: spatio-temporal Poisson

1. time intensity $\lambda > 0$
2. spatial density φ: $\mathbb{P}[\text{demand in } \mathcal{S}] = \int_{\mathcal{S}} \varphi(x)\, dx$
3. inter-arrival times and locations are i.i.d.

Service model:

1. m holonomic vehicles with maximum velocity v
2. vehicles provide a random on-site service
3. on-site service times are i.i.d. (equal on average to $\bar{s}$)
4. demand removed from the system upon on-site service completion

General queueing-theoretical model for DVR 2/2

Performance measure: steady-state system time of demands $\overline{T}$

Problem statement

Solve optimization problem over all causal routing policies π:

$$\inf_{\pi} \overline{T}_{\pi}$$

Relation to standard queueing systems

- DVR model close to M/G/m queue
- **key difference**: service times are **not** i.i.d. in general

Service time correlations in DVR:

- service time = travel time + on-site service
- FCFS policy
- unconditional expected travel time between two consecutive demands ≈ 0.52.
- conditional expected travel time between two consecutive demands > 0.52.

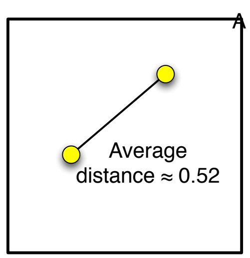

M/G/m methodology is not applicable!

Relation to standard queueing systems

- DVR model close to M/G/m queue
- **key difference**: service times are **not** i.i.d. in general

Service time correlations in DVR:

- service time = travel time + on-site service
- FCFS policy
- unconditional expected travel time between two consecutive demands ≈ 0.52.
- conditional expected travel time between two consecutive demands > 0.52.

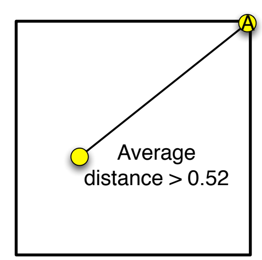

M/G/m methodology is not applicable!

A first look at the problem: stability

- $\lambda \cdot \mathbb{E}[\text{service time}]/m$ fraction of time each vehicle is busy

Necessary condition for stability:
System is **stable** if $\lambda \cdot \mathbb{E}[\text{service time}]/m < 1$.
Since $\bar{s} \leq \mathbb{E}[\text{service time}]$, a weaker necessary condition is:

$$\varrho = \lambda \bar{s}/m < 1$$

Sufficient condition for stability:
Surprisingly, $\varrho < 1$ is also sufficient for stability $\implies$ stability condition is **independent** of the size and shape of $\mathcal{Q}$

Analysis approach

- Lack of i.i.d. property substantially complicates analysis
- General approach:
  1. lower bounds on performance, independent of algorithms,
  2. design of algorithms and upper bound on their performance, possibly in asymptotic regimes (i.e., $\varrho \to 0^+$ and $\varrho \to 1^-$)

Lecture outline

1. Queueing-theoretical model for DVR
2. Lower bounds on performance (m=1)
3. Control policies

D. J. Bertsimas and G. J. van Ryzin. A stochastic and dynamic vehicle routing problem in the Euclidean plane. *Operations Research*, 39:601–615, 1991

D. J. Bertsimas and G. J. van Ryzin. Stochastic and dynamic vehicle routing with general interarrival and service time distributions. *Advances in Applied Probability*, 25:947–978, 1993

Light-load lower bound

Median

- minimizer p^* of

$$p \mapsto \int_{\mathcal{Q}} \|x - p\| \varphi(x) dx = \mathbb{E}_\varphi[\|X - p\|]$$

- best a priori location to reach next demand

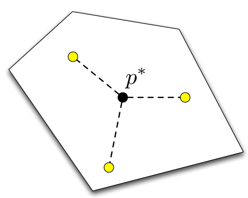

Lower bound (most useful when $\lambda \to 0^+$)

For all policies π: $\overline{T}_\pi \geq \mathbb{E}_\varphi[\|X - p^*\|]/v + \bar{s}$

Proof sketch:

- $\overline{T} = \overline{W}_{\text{travel}} + \overline{W}_{\text{on-site}} + \bar{s}$.
- $\overline{W}_{\text{travel}} \geq \mathbb{E}_\varphi[\|X - p^*\|]/v$

Light-load lower bound

Median

- minimizer p^* of

$$p \mapsto \int_{\mathcal{Q}} \|x - p\| \varphi(x) dx = \mathbb{E}_\varphi[\|X - p\|]$$

- best a priori location to reach next demand

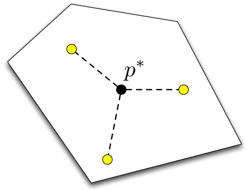

Lower bound (most useful when $\lambda \to 0^+$)

For all policies π: $\overline{T}_\pi \geq \mathbb{E}_\varphi[\|X - p^*\|]/v + \bar{s}$

Proof sketch:

- $\overline{T} = \overline{W}_{\text{travel}} + \overline{W}_{\text{on-site}} + \bar{s}$.
- $\overline{W}_{\text{travel}} \geq \mathbb{E}_\varphi[\|X - p^*\|]/v$

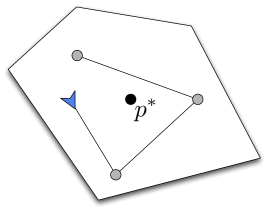

Light-load lower bound

Median

- minimizer p^* of

$$p \mapsto \int_{\mathcal{Q}} \|x - p\| \varphi(x) dx = \mathbb{E}_\varphi[\|X - p\|]$$

- best a priori location to reach next demand

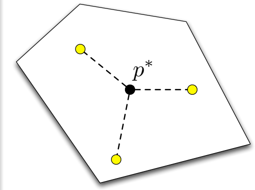

Lower bound (most useful when $\lambda \to 0^+$)

For all policies π: $\overline{T}_\pi \geq \mathbb{E}_\varphi[\|X - p^*\|]/v + \bar{s}$

Proof sketch:

- $\overline{T} = \overline{W}_{\text{travel}} + \overline{W}_{\text{on-site}} + \bar{s}$.
- $\overline{W}_{\text{travel}} \geq \mathbb{E}_\varphi[\|X - p^*\|]/v$

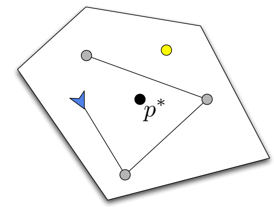

Light-load lower bound

Median

- minimizer p^* of

$$p \mapsto \int_{\mathcal{Q}} \|x - p\| \varphi(x) dx = \mathbb{E}_\varphi[\|X - p\|]$$

- best a priori location to reach next demand

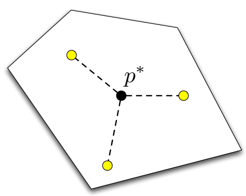

Lower bound (most useful when $\lambda \to 0^+$)

For all policies π: $\overline{T}_\pi \geq \mathbb{E}_\varphi[\|X - p^*\|]/v + \bar{s}$

Proof sketch:

- $\overline{T} = \overline{W}_{\text{travel}} + \overline{W}_{\text{on-site}} + \bar{s}$.
- $\overline{W}_{\text{travel}} \geq \mathbb{E}_\varphi[\|X - p^*\|]/v$

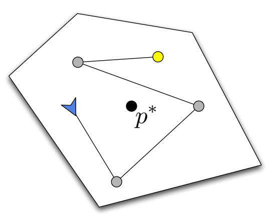

Light-load lower bound

Median

- minimizer p^* of

$$p \mapsto \int_{\mathcal{Q}} \|x - p\| \varphi(x) dx = \mathbb{E}_\varphi[\|X - p\|]$$

- best a priori location to reach next demand

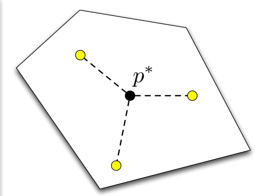

Lower bound (most useful when $\lambda \to 0^+$)

For all policies π: $\overline{T}_\pi \geq \mathbb{E}_\varphi[\|X - p^*\|]/v + \bar{s}$

Proof sketch:

- $\overline{T} = \overline{W}_{\text{travel}} + \overline{W}_{\text{on-site}} + \bar{s}$.
- $\overline{W}_{\text{travel}} \geq \mathbb{E}_\varphi[\|X - p^*\|]/v$

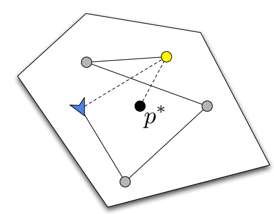

Heavy-load lower bound

Definition (Spatially-biased and -unbiased policies)

A policy π is said to be

1. *spatially unbiased* if system time is independent of demand location
2. *spatially biased* if system time depends on demand location

Heavy-load lower bound

spatially-unbiased policies: $\overline{T}_\pi \geq \frac{\beta_{\mathrm{TSP}}^2}{2} \frac{\lambda \left(\int_Q \varphi^{1/2}(x)dx\right)^2}{v^2 (1-\varrho)^2}$ as $\varrho \to 1^-$

spatially-biased policies: $\overline{T}_\pi \geq \frac{\beta_{\mathrm{TSP}}^2}{2} \frac{\lambda \left(\int_Q \varphi^{2/3}(x)dx\right)^3}{v^2 (1-\varrho)^2}$ as $\varrho \to 1^-$

Heavy-load lower bound

Definition (Spatially-biased and -unbiased policies)

A policy π is said to be

1. *spatially unbiased* if system time is independent of demand location
2. *spatially biased* if system time depends on demand location

Heavy-load lower bound

spatially-unbiased policies: $\overline{T}_\pi \geq \frac{\beta_{\mathrm{TSP}}^2}{2} \frac{\lambda \left(\int_Q \varphi^{1/2}(x)dx\right)^2}{v^2 (1-\varrho)^2}$ as $\varrho \to 1^-$

spatially-biased policies: $\overline{T}_\pi \geq \frac{\beta_{\mathrm{TSP}}^2}{2} \frac{\lambda \left(\int_Q \varphi^{2/3}(x)dx\right)^3}{v^2 (1-\varrho)^2}$ as $\varrho \to 1^-$

Proof sketch (for unbiased policies)

Proof of the lower bound:

- the idea is to use **stability arguments** (which are independent of policies!)
- let $\overline{D}$ be the travel **inter-demand** distance
- one can show that

$$\overline{D} \geq \beta_{\mathrm{TSP}} \frac{\int_Q \varphi^{1/2}(x)dx}{\sqrt{2\,\overline{N}}} \qquad \text{as } \varrho \to 1^-,$$

  with $\overline{N}$ average number of waiting demands
- for stability:

$$\bar{s} + \frac{\overline{D}}{v} \leq \frac{1}{\lambda} \quad \Longrightarrow \quad \bar{s} + \beta_{\mathrm{TSP}} \frac{\int_Q \varphi^{1/2}(x)dx}{v\,\sqrt{2\,\overline{N}}} \leq 1/\lambda$$

- since $\overline{N} = \lambda \overline{W}$ and $\overline{T} = \overline{W} + \bar{s}$ one obtains:

$$\overline{T}^* \geq \frac{\beta_{\mathrm{TSP}}^2}{2} \frac{\lambda \left(\int_Q \varphi^{1/2}(x)dx\right)^2}{v^2 (1-\varrho)^2}$$

Lecture outline

D. J. Bertsimas and G. J. van Ryzin. A stochastic and dynamic vehicle routing problem in the Euclidean plane. *Operations Research*, 39:601–615, 1991

D. J. Bertsimas and G. J. van Ryzin. Stochastic and dynamic vehicle routing with general interarrival and service time distributions. *Advances in Applied Probability*, 25:947–978, 1993

M. Pavone, E. Frazzoli, and F. Bullo. Distributed and adaptive algorithms for vehicle routing in a stochastic and dynamic environment. *IEEE Transactions on Automatic Control*, May 2010. (Submitted, Apr 2009) to appear

An optimal light load policy

Stochastic Queueing Median (SQM)

Compute median p^*. Then:

1: service demands in FCFS order
2: return to p^* after each service is completed

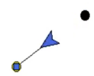

Optimality of SQM policy

$$\lim_{\lambda \to 0^+} \overline{T}_{\mathrm{SQM}}/\overline{T}^* = 1$$

Proof sketch

- As $\lambda \to 0^+$, $\mathbb{P}$ [demand generated when system is empty] $\to 1$
- $\Rightarrow$ all demands generated with the vehicle at p^*
- $\Rightarrow \overline{T}_{\mathrm{SQM}} = \mathbb{E}_\varphi[\|X - p^*\|]/v + \bar{s}$

An optimal light load policy

Stochastic Queueing Median (SQM)

Compute median p^*. Then:

1: service demands in FCFS order
2: return to p^* after each service is completed

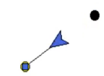

Optimality of SQM policy

$$\lim_{\lambda \to 0^+} \overline{T}_{\mathrm{SQM}}/\overline{T}^* = 1$$

Proof sketch

- As $\lambda \to 0^+$, $\mathbb{P}$ [demand generated when system is empty] $\to 1$
- $\Rightarrow$ all demands generated with the vehicle at p^*
- $\Rightarrow \overline{T}_{\mathrm{SQM}} = \mathbb{E}_\varphi[\|X - p^*\|]/v + \bar{s}$

An optimal spatially-unbiased heavy-load policy

Unbiased TSP (UTSP)

Partition $\mathcal{Q}$ into r subregions $\mathcal{Q}_k$ with $\int_{\mathcal{Q}_k} \varphi(x)dx = 1/r$. Then:

1: within each subregion form sets of size n/r
2: deposit sets in a queue
3: service sets FCFS by following a TSP tour

Optimize over n.

Optimality of UTSP policy

$$\lim_{\varrho \to 1^-} \overline{T}_{\mathrm{UTSP}}(r)/\overline{T}^*_{\mathrm{U}} \leq 1 + 1/r$$

Proof

Proof ($r = 1$)

- idea: reduction to GI/G/1 queue
- jth set viewed as jth customer: arrival and service times are i.i.d.!
- inter-arrival distribution is Erlang of order n
- expected service time is $n\,\bar{s} + \beta_{\mathrm{TSP}}\sqrt{n}\int_{\mathcal{Q}} \varphi^{1/2}(x)dx/v$
- standard results give upper bound on the wait in queue for a set
- then easy to find upper bound for individual demands

Comments

Relation with non-spatial queueing systems:

- wait time grows as $(1-\varrho)^{-2}$ instead of $(1-\varrho)^{-1}$!
- DVR problems are fundamentally different from traditional queueing systems (techniques, results, etc.)

Analysis techniques:

- for light load: **locational optimization**
- for heavy load: **reduction to classic queueing systems** or **control-theoretical methods**

Biased/unbiased:

- biased service provides strict reduction of optimal system time for any non-uniform φ

Adaptivity

SQM policy not adaptive:

- SQM unstable as $\varrho \to 1^-$
- intuition: average per-demand travel $\overline{D}$ is **fixed**
- but stability condition implies $\overline{D} < (1-\varrho)/\lambda$!

UTSP and BTSP policies not adaptive:

- for stability of the queue of sets:

$$\frac{\lambda}{n}\left(n\,\bar{s} + \beta_{\mathrm{TSP}}\sqrt{n}\int_{\mathcal{Q}} \varphi^{1/2}(x)dx/v\right) < 1$$

- then one should **a priori** select:

$$n > \lambda^2 \beta_{\mathrm{TSP}}^2 \left[\int_{\mathcal{Q}} \varphi^{1/2}(x)\,dx\right]^2 / (v^2\,(1-\varrho)^2)$$

- $\Rightarrow$ wrong selection of n might lead to instability or unacceptable deterioration in performance

Divide & Conquer policy

Divide & Conquer (DC)

Partition $\mathcal{Q}$ into r subregions $\mathcal{Q}_k$ with $\int_{\mathcal{Q}_k} \varphi(x)dx = 1/r$. Then:

1: while no customers, move to empirical median $\tilde{p}^*$

2: while customers waiting

1. move to subregion $\mathcal{Q}_k$
2. service all demands in $\mathcal{Q}_k$ by following a TSP tour
3. $k \leftarrow k+1$ (modulo r)

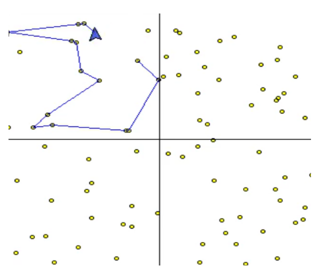

DC policy (with $r \to +\infty$)

Implementation:

- NP-hard computation, but effective heuristics

Adaptation: the policy does not require knowledge of

1. vehicle velocity v, environment $\mathcal{Q}$
2. arrival rate λ
3. expected on-site service $\bar{s}$

Performance:

1. in light load, delay is optimal
2. in heavy load, delay is optimal
3. stable in any load condition

optimal and adaptive
very little known outside of asymptotic regimes

Proof (r=1)

Light load:

- $\tilde{p}^* \to p^*$ and recovers SQM

Heavy load:

- no well-defined notion of "jth customer"
- focus on dynamical system

$$\mathbb{E}[n_{i+1}] \leq \lambda \, \mathbb{E}\Big[\sum_{q=1}^{n_i} s_q + \mathrm{TSP}(n_i)\Big]$$
$$\leq \lambda \left(\bar{s} \; \mathbb{E}[n_i] + \beta_{\mathrm{TSP}} \int_{\mathcal{Q}} \varphi^{1/2}(x) dx \sqrt{\mathbb{E}[n_i]}/v\right)$$

- upper bound trajectories with the trajectories of **virtual** dynamical system

$$z_{i+1} = \varrho \, z_i + (\lambda/v) \, \beta_{\mathrm{TSP}} \int_{\mathcal{Q}} \varphi^{1/2}(x) dx \sqrt{z_i}$$

- $\overline{T}_{\mathrm{DC}} \leq \lim_{i \to +\infty} z_i/\lambda$

Receding-Horizon policy

Receding-Horizon (RH)

For $\eta \in (0, 1]$, single agent performs:

```
1: while no customers, move to empirical median p̃*
2: while customers waiting
   1 compute TSP tour through current demands
   2 service η-fraction of path
```

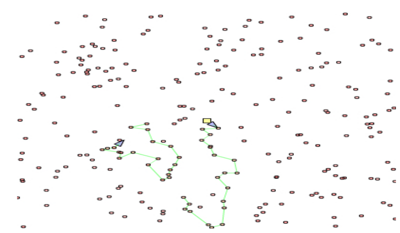

RH policy

Implementation:

- NP-hard computation, but effective heuristics

Adaptation: the policy does not require knowledge of

1. vehicle velocity v, environment $\mathcal{Q}$
2. arrival rate λ and spatial density function φ
3. expected on-site service $\bar{s}$

Performance:

1. in light load, delay is optimal
2. in heavy load, delay is within a multiplicative factor from optimal
3. multiplicative factor depends upon φ and is conjectured to equal 2

adaptive to all parameters

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