Dynamic Vehicle Routing for Robotic Networks
Lecture #1: Introduction

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

1. Acknowledgements
2. Autonomy and Networking Technologies
3. Prototypical DVR problem
4. Literature review
5. Contributions
6. Comparison with alternative approaches
   - Re-optimization
   - Online algorithms
7. Workshop Structure and Schedule

Acknowledgements

Funded in part by
AFOSR grant no. FA 8650-07-2-3744 (Michigan/AFRL Collaborative Center on Control Sciences), ARO MURI "Swarms" W911NF-05-1-0219, ARO award DAAD19-03-D-0004 (Institute for Collaborative Biotechnologies), NSF awards #0705451 and #0705453, ONR award N00014-07-1-0721.

Collaborators
Alessandro Arsie (U. of Toledo), Shaunak D. Bopardikar (UCSB), Ruggero Carli (UCSB/Padova), Jorge Cortés (UCSD), Joey W. Durham (UCSB), John J. Enright (Kiva Systems), Paolo Frasca (Roma), Anurag Ganguli (UtopiaCompression), João P. Hespanha (UCSB) Sonia Martínez (UCSD), Karl Obermeyer (UCSB), and Sara Susca (Honeywell).

Autonomy and Networking Technologies

Individual members in the group can
- **sense** its immediate environment
- **communicate** with others
- **process** the information gathered
- **take a local action** in response

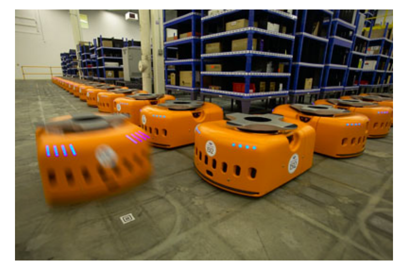

Lecture outline

1. Acknowledgements
2. Autonomy and Networking Technologies
3. **Prototypical DVR problem**
4. Literature review
5. Contributions
6. Comparison with alternative approaches
   - Re-optimization
   - Online algorithms
7. Workshop Structure and Schedule

Prototypical Dynamic Vehicle Routing Problem

Given:

- a group of vehicles, and
- a set of service demands

Objective:
provide service in minimum time
service = take a picture at location

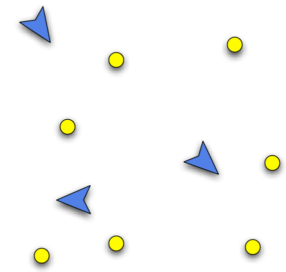

Vehicle routing (All info known ahead of time, Dantzig '59)

Determine a set of paths that allow vehicles to service the demands

Dynamic vehicle routing (New info in real time, Psaraftis '88)

- New demands arise in real-time
- Existing demands evolve over time

Prototypical Dynamic Vehicle Routing Problem

Given:

- a group of vehicles, and
- a set of service demands

Objective:
provide service in minimum time
service = take a picture at location

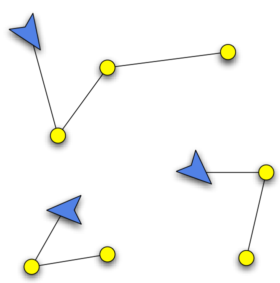

Vehicle routing (All info known ahead of time, Dantzig '59)

Determine a set of paths that allow vehicles to service the demands

Dynamic vehicle routing (New info in real time, Psaraftis '88)

- New demands arise in real-time
- Existing demands evolve over time

Prototypical Dynamic Vehicle Routing Problem

Given:

- a group of vehicles, and
- a set of service demands

Objective:
provide service in minimum time
service = take a picture at location

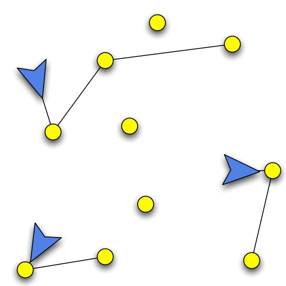

Vehicle routing (All info known ahead of time, Dantzig '59)

Determine a set of paths that allow vehicles to service the demands

Dynamic vehicle routing (New info in real time, Psaraftis '88)

- New demands arise in real-time
- Existing demands evolve over time

Prototypical Dynamic Vehicle Routing Problem

Given:

- a group of vehicles, and
- a set of service demands

Objective:
provide service in minimum time
service = take a picture at location

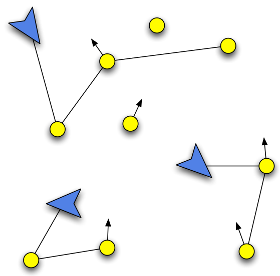

Vehicle routing (All info known ahead of time, Dantzig '59)

Determine a set of paths that allow vehicles to service the demands

Dynamic vehicle routing (New info in real time, Psaraftis '88)

- New demands arise in real-time
- Existing demands evolve over time

Light and heavy load regimes

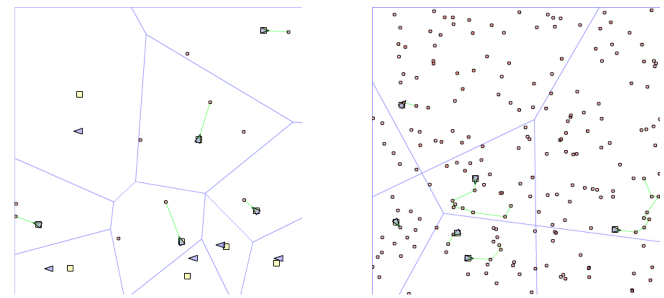

Lecture outline

1. Acknowledgements
2. Autonomy and Networking Technologies
3. Prototypical DVR problem
4. Literature review
5. Contributions
6. Comparison with alternative approaches
   - Re-optimization
   - Online algorithms
7. Workshop Structure and Schedule

From coordination and static routing to Dynamic Vehicle Routing

Simple coordination problems arise in static environments

1. motion coordination: rendezvous, deployment, flocking
2. task allocation, target assignment
3. static vehicle routing (P. Toth and D. Vigo '01)

Routing policies vs planning algorithms

dynamic, stochastic and adversarial events take place

1. design **policies** (in contrast to pre-planned routes or motion planning algorithms) to specify how to react to events
2. dynamic demands add **queueing phenomena** to the combinatorial nature of vehicle routing

Literature on DVR and queueing for robotic networks

- Shortest path through randomly-generated and worst-case points (Beardwood, Halton and Hammersly, 1959 — Steele, 1990)
- Traveling salesman problem solvers (Lin, Kernighan, 1973)
- DVR formulation on a graph (Psaraftis, 1988)
- DVR on Euclidean plane (Bertsimas and Van Ryzin, 1990–1993)
- Unified receding-horizon policy (Papastavrou, 1996)

Recent developments in DVR for robotic networks:

- Adaptation and decentralization
- Vehicles with dynamics, nonholonomic vehicles, Dubins UAVs
- Pickup & delivery tasks
- Heterogeneous vehicles and team forming
- Distinct-priority and impatient demands
- Moving demands

Lecture outline

1. Acknowledgements
2. Autonomy and Networking Technologies
3. Prototypical DVR problem
4. Literature review
5. Contributions
6. Comparison with alternative approaches
   - Re-optimization
   - Online algorithms
7. Workshop Structure and Schedule

Contributions of our recent works

Comprehensive framework for DVR in robotic systems

1. adaptive DVR policies for single vehicles in light and heavy load
2. cooperative DVR policies via partitioning
3. scalable distributed partitioning policies under a variety of communication/interaction scenarios
4. (models, algorithms and analysis of) service vehicles with dynamics & stochastic and combinatorics of nonholonomic Dubins vehicles performing Traveling Salesman Problems and DVR tasks
5. (models, algorithms and analysis of) service vehicles with time constraints and heterogeneous priorities
6. (models, algorithms and analysis of) demands requiring service by multiple heterogeneous vehicles simultaneously.

Bibliography on DVR and queueing for robotic networks

1. K. Savla, E. Frazzoli, and F. Bullo. Traveling Salesperson Problems for the Dubins vehicle. *IEEE Transactions on Automatic Control*, 53(6):1378–1391, 2008
2. K. Savla, F. Bullo, and E. Frazzoli. Traveling Salesperson Problems for a double integrator. *IEEE Transactions on Automatic Control*, 54(4):788–793, 2009
3. J. J. Enright, K. Savla, E. Frazzoli, and F. Bullo. Stochastic and dynamic routing problems for multiple UAVs. *AIAA Journal of Guidance, Control, and Dynamics*, 34(4):1152–1166, 2009
4. 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
5. 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
6. A. Arsie, K. Savla, and E. Frazzoli. Efficient routing algorithms for multiple vehicles with no explicit communications. *IEEE Transactions on Automatic Control*, 54(10):2302–2317, 2009
7. 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
8. M. Pavone, K. Savla, and E. Frazzoli. Sharing the load. *IEEE Robotics and Automation Magazine*, 16(2):52–61, 2009
9. M. Pavone, A. Arsie, E. Frazzoli, and F. Bullo. Equitable partitioning policies for mobile robotic networks. *IEEE Transactions on Automatic Control*, 2010. (Submitted Dec 2008 and Aug 2009) to appear
10. 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
11. S. D. Bopardikar, S. L. Smith, F. Bullo, and J. P. Hespanha. Dynamic vehicle routing for translating demands: Stability analysis and receding-horizon policies. *IEEE Transactions on Automatic Control*, 55(11), 2010. (Submitted, Mar 2009) to appear
12. S. D. Bopardikar, S. L. Smith, and F. Bullo. On vehicle placement to intercept moving targets. *Automatica*, March 2010. Submitted
13. F. Bullo, E. Frazzoli, M. Pavone, K. Savla, and S. L. Smith. Dynamic vehicle routing for robotic systems. *Proceedings of the IEEE*, May 2010. Submitted
14. H. A. Waisanen, D. Shah, and M. A. Dahleh. A dynamic pickup and delivery problem in mobile networks under information constraints. *IEEE Transactions on Automatic Control*, 53(6):1419–1433, 2008
15. R. Szechtman, M. Kress, K. Lin, and D. Cfir. Models of sensor operations for border surveillance. *Naval Research Logistics*, 55(1):27–41, 2008
16. S. Itani. *Dynamic Systems and Subadditive Functionals*. PhD thesis, Massachusetts Institute of Technology, 2009

Lecture outline

1. Acknowledgements
2. Autonomy and Networking Technologies
3. Prototypical DVR problem
4. Literature review
5. Contributions
6. Comparison with alternative approaches
   - Re-optimization
   - Online algorithms
7. Workshop Structure and Schedule

Plain-vanilla re-optimization?

Example: DVR on segment

- Objective: minimize average waiting time
- Strategy: re-optimize at each event

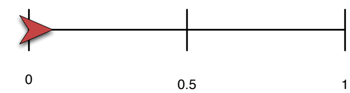

1. For adversarial target generation, vehicle travels forever without ever servicing any request $\Longrightarrow$ **unstable queue of outstanding requests**
2. Even if queue remains bounded, what about performance? how far from the optimal?

Plain-vanilla re-optimization?

Example: DVR on segment

- Objective: minimize average waiting time
- Strategy: re-optimize at each event

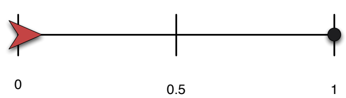

1. For adversarial target generation, vehicle travels forever without ever servicing any request $\Longrightarrow$ **unstable queue of outstanding requests**
2. Even if queue remains bounded, what about performance? how far from the optimal?

Plain-vanilla re-optimization?

Example: DVR on segment

- Objective: minimize average waiting time
- Strategy: re-optimize at each event

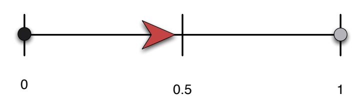

1. For adversarial target generation, vehicle travels forever without ever servicing any request $\Longrightarrow$ **unstable queue of outstanding requests**
2. Even if queue remains bounded, what about performance? how far from the optimal?

Plain-vanilla re-optimization?

Example: DVR on segment

- Objective: minimize average waiting time
- Strategy: re-optimize at each event

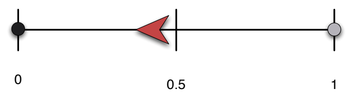

1. For adversarial target generation, vehicle travels forever without ever servicing any request $\Longrightarrow$ **unstable queue of outstanding requests**
2. Even if queue remains bounded, what about performance? how far from the optimal?

Plain-vanilla re-optimization?

Example: DVR on segment

- Objective: minimize average waiting time
- Strategy: re-optimize at each event

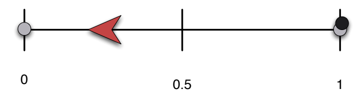

1. For adversarial target generation, vehicle travels forever without ever servicing any request ⟹ unstable queue of outstanding requests
2. Even if queue remains bounded, what about performance? how far from the optimal?

FB, EF, MP, KS, SLS (UCSB, MIT) Dynamic Vehicle Routing (Lecture 1/8) 29jun10 @ Baltimore, ACC 15 / 18

Plain-vanilla re-optimization?

Example: DVR on segment

- Objective: minimize average waiting time
- Strategy: re-optimize at each event

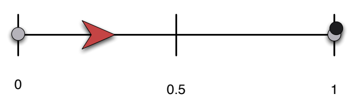

1. For adversarial target generation, vehicle travels forever without ever servicing any request ⟹ unstable queue of outstanding requests
2. Even if queue remains bounded, what about performance? how far from the optimal?

FB, EF, MP, KS, SLS (UCSB, MIT) Dynamic Vehicle Routing (Lecture 1/8) 29jun10 @ Baltimore, ACC 15 / 18

Plain-vanilla re-optimization?

Example: DVR on segment

- Objective: minimize average waiting time
- Strategy: re-optimize at each event

1. For adversarial target generation, vehicle travels forever without ever servicing any request ⟹ unstable queue of outstanding requests
2. Even if queue remains bounded, what about performance? how far from the optimal?

FB, EF, MP, KS, SLS (UCSB, MIT) Dynamic Vehicle Routing (Lecture 1/8) 29jun10 @ Baltimore, ACC 15 / 18

Online algorithms?

Online algorithms (Jaillet and M. R. Wagner '06)

- online algorithm operates based on input information up to the current time
- online algorithm is (worst-case) r-competitive if

$$\text{Cost}_{\text{online}}(I) \leq r\text{Cost}_{\text{optimal offline}}(I), \qquad \forall \text{ problem instances } I.$$

Disadvantages

1. cumulative cost
2. worst-case analysis
3. not possible to include *a-priori* information (e.g., arrival rate)
4. not as clear what competitive ratio means
5. so far, only few simple DVR problems admit online algorithms

Lecture outline

1. Acknowledgements
2. Autonomy and Networking Technologies
3. Prototypical DVR problem
4. Literature review
5. Contributions
6. Comparison with alternative approaches
   - Re-optimization
   - Online algorithms
7. Workshop Structure and Schedule

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