**Workshop for the 2010 American Control Conference**
**Baltimore, Maryland, USA, June 30 - July 2, 2010**

# Dynamic Vehicle Routing for Robotic Systems


## Organizers and Lecturers

**Francesco Bullo**
Affiliation: Mechanical Engineering Department, University of California at Santa Barbara
Address: ME, 2338 Engineering Bldg II, Santa Barbara, CA 93106-5070, USA
Email: `bullo@engr.ucsb.edu`, http://motion.me.ucsb.edu, Ph 805.893.5169, Fax 805.893.8651

**Emilio Frazzoli**
Affiliation: Laboratory for Information and Decision Systems, and Aeronautics and Astronautics Department, Massachusetts Institute of Technology
Address: 77 Massachusetts Avenue, Room 33-332, Cambridge, MA 02139, USA
Email: `frazzoli@mit.edu`, http://ares.lids.mit.edu, Ph 617.253.1991

**Marco Pavone**
Affiliation: Laboratory for Information and Decision Systems, and Aeronautics and Astronautics Department, Massachusetts Institute of Technology
Address: 77 Massachusetts Avenue, Room 32-D712, Cambridge, MA 02139, USA
Email: `pavone@mit.edu`, http://mit.edu/pavone/www/, Ph 617.501.7485

**Ketan Savla**
Affiliation: Laboratory for Information and Decision Systems, Massachusetts Institute of Technology
Address: 77 Massachusetts Avenue, Room 32-D716, Cambridge, MA 02139, USA
Email: `ksavla@mit.edu`, http://mit.edu/ksavla/www/, Ph 617.324.0095

**Stephen L. Smith**
Affiliation: Computer Science & Artificial Intelligence Laboratory, Massachusetts Institute of Technology
Address: 32 Vassar Street, 32-376, Cambridge, MA 02139, USA
Email: `slsmith@mit.edu`, http://mit.edu/slsmith/www/, Ph 617.253.6532, Fax 617.253.6849


# Workshop overview

The workshop presents a joint algorithmic and queueing approach to the design of cooperative control and task allocation strategies for networks of uninhabited vehicles (UVs). The key novelty is the introduction of *stochastic, combinatorial and queueing aspects* in the context of distributed coordination of multi-agent networks. The work builds upon our recent breakthroughs in the design of *dynamic vehicle routing (DVR)* algorithms; these algorithms lead vehicles to complete dynamically-generated tasks with random locations and characteristics.

We present recent work on coordination, planning and routing algorithms for the efficient discovery and timely servicing of tasks that are not known a priori. As in queueing theory, task arrival is modeled as a stochastic process, and queueing-style algorithms are required to enable UVs to search, identify, allocate, prioritize, plan paths, and form teams. DVR algorithms are typically based on a combination of receding-horizon resource allocation, distributed optimization, combinatorics and control.

# Scientific rationale: static vs dynamic vehicle routing

### Static vehicle routing

Over the last few years the controls community, in collaboration with scientists in robotics and network science, has made great strides in developing a formal understanding of multi-agent networks such as networks of autonomous vehicles. A useful characterization is now available of numerous basic coordination tasks such as consensus, rendezvous, deployment, and flocking; our recent monograph (Bullo et al., 2009) contains a coherent introduction to robotic coordination and (Tanner et al., 2007; Olfati-Saber et al., 2007; Krick et al., 2009) discuss other relevant approaches. Furthermore, numerical optimization methods and other algorithmic approaches have had a good measure of success in dealing with more DoD-relevant coordination objectives such as the task allocation problems; for example see (Moore and Passino, 2007; Smith and Bullo, 2009a; Alighanbari and How, 2008; Beard et al., 2002). In these papers, the underlying mathematical model is *static* and fits within the framework of the Vehicle Routing Problem (see (Toth and Vigo, 2001) for a thorough introduction). The vehicle routing problem is described as follows: (a) a team of vehicles is required to service a set of demands; (b) each demand requires a certain amount of on-site service; (c) the goal is to compute a set of routes that optimizes the cost of servicing the demands (according to some quality of service metric). The vehicle routing problem is *static* in the sense that vehicle routes are computed assuming that no new demands arrive. In other words, most of the available literature focuses on static environments and does not properly account for scenarios in which dynamic, stochastic and adversarial events take place.

### From static to dynamic vehicle routing

The problem of planning routes through service demands that arrive or are discovered during a mission execution is known as the "stochastic and dynamic vehicle routing problem" (abbreviated as the DVR problem in the operations research literature). There are two key differences between static and dynamic vehicle routing problems. First, planning algorithms should actually provide *policies* (in contrast to pre-planned routes) that prescribe how the routes should evolve as a function of those inputs that evolve in real-time. Second, dynamic demands (i.e., demands that vary over time) add *queueing phenomena* to the combinatorial nature of vehicle routing.

The simplest DVR problem (a vehicle moves along straight lines and visits targets whose time of arrival, location and on-site service are stochastic) was investigated by Bertsimas and van Ryzin (1991, 1993); Papastavrou (1996); see also the earlier work (Psaraftis, 1980) where the vehicle and targets locations belong to a graph.

### Our recent work on multi-vehicle and robotic networks

Starting with our initial works (Savla et al., 2005), we have recently studied several variations of the basic DVR problem, including (in rough chronological order) single vehicles with nonholonomic constraints (Savla et al., 2008a), groups of vehicles with nonholonomic constraints (Enright et al., 2009), decentralized and adaptive policies (Pavone et al., 2010), group of vehicles with no communication (Arsie et al., 2009) or limited sensing range (Enright and Frazzoli, 2006), tasks that require teams (Smith and Bullo, 2009b), impatient tasks (Pavone et al., 2009), tasks with distinct importance (Smith et al., 2010), translating tasks (Bopardikar et al., 2010), and human-in-the-loop routing policies (Savla et al., 2008b). These recent developments are the focus of the workshop.

We discuss dynamic vehicle routing problems in their full generality. We describe an algorithmic queueing theory for robotic networks by integrating dynamics, combinatorial optimization, teaming, and distributed algorithms. For each scenario, we present a three-step approach:

1. the analysis of the problem structure (by this we mean for example (a) a dimensional analysis to focus on dimension-less parameters, (b) the identification of intrinsic regimes with distinct problem characteristics, and (c) the study of phase transitions between intrinsic regimes);

2. the establishment of fundamental limits on performance, independent of algorithms;

3. the design of algorithms that are either optimal or constant-factor away from optimal, that are adaptive and require no a-priori information about the environment and the problem data, and that require limited information exchanges; and

## Intended workshop audience

The intended audience of the workshop are graduate students and researchers in control robotics and queueing theory from Computer Science, Electrical Engineering, Mechanical Engineering, and Aerospace Engineering. The proposal of this workshop is timely. The last years have seen a thriving research activity on vehicle routing and robotic autonomy. This interest is motivated by the growing possibilities enabled by robotic networks in the monitoring of natural phenomena and the enhancement of human capabilities in hazardous and unknown environments. The workshop presents a coherent introduction to the application of robotic systems to a variety of spatially distributed tasks.

## Workshop outline

The workshop will consist of the following lectures.

**Lecture #1: Introduction to DVR**
(1) problem statement and literature review; (2) motivation for queueing theory, applications to robotic systems and to distributed real-time autonomy and decision making and comparison with online algorithms from operations research

**Lecture #2: Prelims: graphs, TSPs, and queues**
elements of (1) graph theory, (2) TSP and combinatorial problems, (3) queueing theory

**Lecture #3: The single-vehicle DVR problem**
(0) definition of median, (1) classic results (Bertsimas and van Ryzin 91 and Papastavrou 96), (2) adaptive policy for light and heavy load, and (3) heuristics improvements

**Lecture #4: The multi-vehicle DVR problem**
(1) facility location and Lloyd algorithm, (2) classic results (Bertsimas and van Ryzin 93), (3) motivation/theorems for partitioning

**Lecture #5: Extensions to vehicle networks**
(1) intro to communication models, (2) distributed partitioning policies via power diagrams, gossip, and MacQueen strategies

**Lecture #6: Extensions to different demand models**
models, algorithms and analysis of demands with time constraints and heterogeneous priorities

**Lecture #7: Extensions to different vehicle models**
(1) different vehicle models with differential constraints, e.g., Dubins, double integrator, and others, (2) Dubins minimum paths, (3) models, algorithms and analysis of service vehicles with nonholonomic Dubins dynamics

**Lecture #8: Extensions to different task models**
models, algorithms and analysis of demands requiring service by multiple heterogeneous vehicles simultaneously

## Schedule and Speakers

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

## Brief lecturers Bios

**Francesco Bullo** received the Laurea degree ”summa cum laude” in Electrical Engineering from the University of Padova, Italy, in 1994, and the Ph.D. degree in Control and Dynamical Systems from the California Institute of Technology in 1999. From 1998 to 2004, he was an Assistant Professor with the Coordinated Science Laboratory at the University of Illinois at Urbana-Champaign. He is currently a Professor with the Mechanical Engineering Department at the University of California, Santa Barbara. His research interests include motion coordination for multi-agent networks, motion planning for autonomous vehicles, and geometric control of mechanical systems. He is a recipient of the 2003 ONR Young Investigator Award. He is the coauthor, with Andrew D. Lewis, of the book ”Geometric Control of Mechanical Systems” (Springer, 2004, 0-387-22195-6) and, with Jorge Cortes and Sonia Martinez, of the book ”Distributed Control of Robotic Networks” (Princeton, 2009, 978-0-691-14195-4). He has published more than 150 papers in international journals, books, and refereed conferences. He served or is serving on the Editorial Board of the “SIAM Journal of Control and Optimization” and “IEEE Transactions on Automatic Control.”

**Emilio Frazzoli** is an Associate Professor of Aeronautics and Astronautics with the Laboratory for Information and Decision Systems at the Massachusetts Institute of Technology. He received a Laurea degree in Aerospace Engineering from the University of Rome, ”Sapienza” , Italy, in 1994, and a Ph. D. degree in Navigation and Control Systems from the Department of Aeronautics and Astronautics of the Massachusetts Institute of Technology, in 2001. Between 1994 and 1997 he worked as an officer in the Italian Navy, and as a spacecraft dynamics specialist for the European Space Agency Operations Centre (ESOC) in Darmstadt, Germany, and Telespazio, in Rome, Italy. From 2001 to 2004 he was an Assistant Professor of Aerospace Engineering at the University of Illinois at Urbana-Champaign. From 2004 to 2006 he was an Assistant Professor of Mechanical and Aerospace Engineering at the University of California, Los Angeles. He is an Associate Fellow of the American Institute of Aeronautics and Astronautics and a Senior Member of the Institute for Electrical and Electronics Engineers. He was the recipient of a NSF CAREER award in 2002. He is currently serving on the Editorial Board of the AIAA Journal of Guidance, Control, and Dynamics.

**Marco Pavone** is a Ph.D. Candidate in the Laboratory for Information and Decision Systems within the Department of Aeronautics and Astronautics at the Massachusetts Institute of Technology. Before joining MIT he was a visiting scholar at the University of California, Los Angeles, and he worked as an Analyst for Accenture Consulting. He received a Laurea degree in Systems and Control Engineering from the University of Catania, Italy in 2004, and earned a Diploma Engineer degree from Scuola Superiore di Catania, Italy in 2005. His research interests include algorithmic and computational approaches to the design and development of decision and control architectures for complex networked and autonomous systems, and bio-inspired robotics.

**Ketan Savla** is a Research Scientist in the Laboratory for Information and Decision Systems at the Massachusetts Institute of Technology. He received the B. Tech. degree in Mechanical Engineering from the Indian Institute of Technology, Bombay in 2003, the M.S. degree in Mechanical Engineering from the University of Illinois at Urbana-Champaign in 2004, the M.A. degree in Applied Mathematics and the Ph.D. degree in Electrical Engineering, both from the University of California at Santa Barbara in 2007. His research interests span mobile cyber-physical systems, human-in-loop control systems, intelligent transportation networks and learning-based control algorithms for mobile robotic systems. His CDC-ECC’05 paper was a finalist for the Best Student Paper Award.

**Stephen L. Smith** is a Postdoctoral Associate with the Computer Science and Artificial Intelligence Laboratory at the Massachusetts Institute of Technology. He received the B.Sc. degree in Engineering Physics from Queen’s University, Canada in 2003, the M.A.Sc. degree in Electrical Engineering from the University of Toronto, Canada in 2005, and the Ph.D. degree in Mechanical Engineering from the University of California, Santa Barbara in 2009. He has also held visiting research positions in the Laboratory for Information and Decision Systems at the Massachusetts Institute of Technology and in the GRASP Laboratory at the University of Pennsylvania. His research focuses on vehicle routing, task allocation, and distributed control for autonomous systems. He was a finalist for the Best Student Paper Award at the 2007 IEEE Conference on Decision and Control.

# References


M. Alighanbari and J. P. How. A robust approach to the UAV task assignment problem. *International Journal on Robust and Nonlinear Control*, 18(2):118–134, 2008.

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.

R. W. Beard, T. W. McLain, M. A. Goodrich, and E. P. Anderson. Coordinated target assignment and intercept for unmanned air vehicles. *IEEE Transactions on Robotics and Automation*, 18(6):911–922, 2002.

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 in the Euclidean plane with multiple capacitated vehicles. *Operations Research*, 41(1):60–76, 1993.

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.

F. Bullo, J. Cortés, and S. Martínez. *Distributed Control of Robotic Networks*. Applied Mathematics Series. Princeton University Press, 2009. ISBN 978-0-691-14195-4. Available at http://www.coordinationbook.info.

J. J. Enright and E. Frazzoli. Cooperative UAV routing with limited sensor range. In *AIAA Conf. on Guidance, Navigation and Control*, Keystone, CO, Aug. 2006.

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.

L. Krick, M. E. Broucke, and B. Francis. Stabilization of infinitesimally rigid formations of multi-robot networks. *International Journal of Control*, 82(3):423–439, 2009.

B. J. Moore and K. M. Passino. Distributed task assignment for mobile agents. *IEEE Transactions on Automatic Control*, 52(4):749–753, 2007.

R. Olfati-Saber, J. A. Fax, and R. M. Murray. Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1):215–233, 2007.

J. D. Papastavrou. A stochastic and dynamic routing policy using branching processes with state dependent immigration. *European Journal of Operational Research*, 95:167–177, 1996.

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.

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.

H. N. Psaraftis. Dynamic programming solution to the single vehicle many-to-many immediate request dial-a-ride problem. *Transportation Science*, 14(2):130–154, 1980.

K. Savla, E. Frazzoli, and F. Bullo. On the point-to-point and traveling salesperson problems for Dubins' vehicle. In *American Control Conference*, pages 786–791, Portland, OR, June 2005.

K. Savla, E. Frazzoli, and F. Bullo. Traveling Salesperson Problems for the Dubins vehicle. *IEEE Transactions on Automatic Control*, 53(6):1378–1391, 2008a.

K. Savla, T. Temple, and E. Frazzoli. Human-in-the-loop vehicle routing policies for dynamic environments. In *IEEE Conf. on Decision and Control*, pages 1145–1150, Cancún, México, Dec. 2008b.

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

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, 2009b.

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.

H. G. Tanner, A. Jadbabaie, and G. J. Pappas. Flocking in fixed and switching networks. *IEEE Transactions on Automatic Control*, 52(5):863–868, 2007.

P. Toth and D. Vigo, editors. *The Vehicle Routing Problem*. Monographs on Discrete Mathematics and Applications. SIAM, 2001. ISBN 0898715792.