Dynamic Routing and Robotic Coordination for Oceanographic Adaptive Sampling

Francesco Bullo
Subhash Suri
Gaurav Sukhatme

UCSB
USC Viterbi School of Engineering

Objective

Our approach is a joint theoretical and experimental effort to tackle fundamental questions in coordination for robotic networks: how are tasks partitioned among robots? in what order are they to be performed? along which routes do individual robots move? Specifically, we study resource allocation, vehicle routing, and path planning problems and their application to oceanography via our USC Networked Aquatic Platforms.

Applications

Ocean sampling: underwater gliders to collect data in the Southern California Bight.
Patrolling: cameras and mobile robots for persistentl patrolling.
Logistics: autonomous robots to provide time-aware and efficient delivery of goods.

Year 1: One-to-Base Partitioning

We have explored how to partition space into territories to optimize expected distances to tasks.

One-to-base Cost Function:

$$\mathcal{H}_{\max}(c,P) = \sum_{k\in Q} \max_i \{\mathrm{dist}(c_i,k) \mid k \in P_i\}$$

When robot i talks to base:

1: Update robot i's centroid
2: Transmit local copy of P_i to robot i
3: **for** every other robot j
4: Add vertices from P_i to P_j that are closer to j
5: Remove vertices from P_j that are closer to i

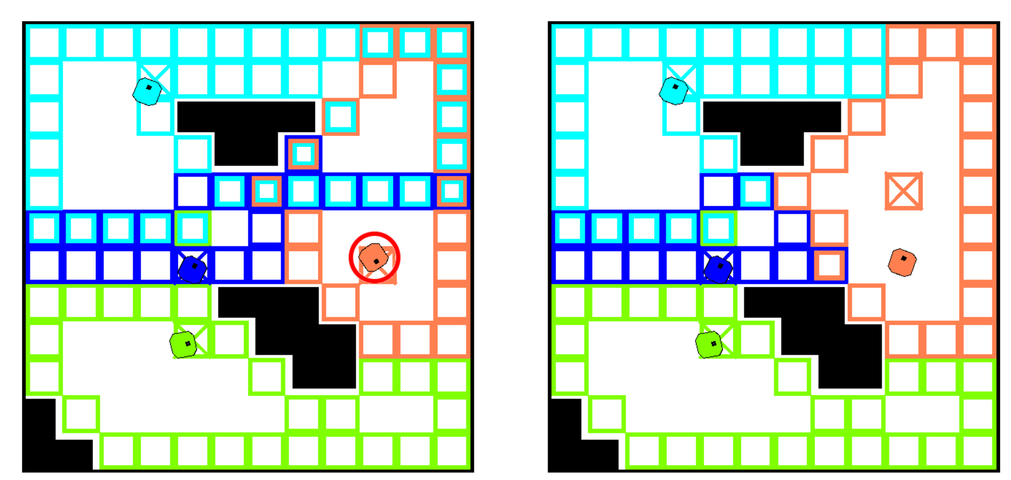

J. W. Durham, R. Carli, P. Frasca, and F. Bullo, "Dynamic Coverage Control with Asynchronous One-to-Base-Station Communication", IEEE Conf. Decision and Control, Dec 2011

Year 2: Cooperative Patrolling

We have explored how to surveil an environment with a team of autonomous robots.

Refresh Time Cost Function: longest weighted time between two consecutive visits of viewpoints

$$\mathrm{RT}(X) = \max\{\phi_\alpha(t_a(\alpha,t_d) - t_d)\}$$

where $t_a(\alpha, t_d)$ is the earliest arrival time by any robot at viewpoint α after departure at time t_d.

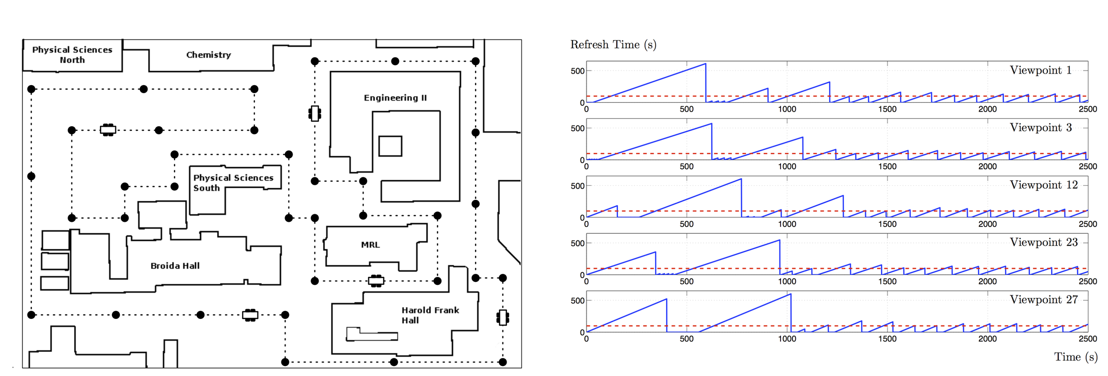

F. Pasqualetti, J. W. Durham, and F. Bullo, "Cooperative Patrolling via Weighted Tours: Performance Analysis and Distributed Algorithms", IEEE Transaction on Robotics, 2012

Route Planning over Stochastic Regions

We consider the *stochastic* version of the classical traveling salesman problem with neighborhoods (TSPN).

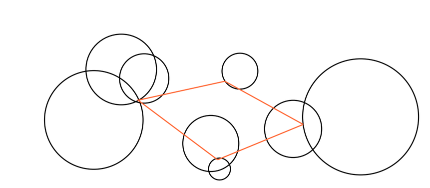

The tour must visit a set of (possibly intersecting) disks.

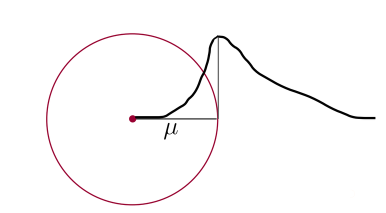

The radius of each disk is an ***i.i.d.*** random variable drawn from a known probability distribution with mean μ.

In this setting, we look at:

- Obtaining an **offline** *order* on the disks, and a *strategy* to move between them, such that the *expected* length of the resulting tour is competitive to the expectation of the optimal.
- Obtaining a tour in the *online* setting, where the radius of each disk is revealed only upon reaching its boundary.

Hardness and Approximation Results

- The stochastic TSPN problems is ***NP***-Hard as a generalization of the classical TSP problem.
- We propose a polynomial-time algorithm that achieves a factor $O(\log\log n)$ approximation of the *expected length of the optimal tour* in the offline case.
- For the *online* version of the problem, we achieve an approximation ratio of $O(\log n)$.
- When the input disks with their *mean radii* are (nearly) disjoint, we achieve a constant-factor approximation even for the online case.

This problem arises in data gathering applications, where a *data mule* needs to collect data from n geographically distributed wireless sensor nodes, whose communication range r is a random variable subject to environmental factors.

P. Kamousi and S. Suri, "Euclidean Traveling Salesman Tours through Stochastic Neighborhoods", *Symposium on Discrete Algorithms (SODA 2013) (submitted).*

Minimum Risk Path Planning with Ocean Current Predictions

- Problem: Plan paths to minimize collision risk for AUVs in the presence of ocean currents
- Utilize current predictions from Ocean Models
- Account for uncertainty in current predictions

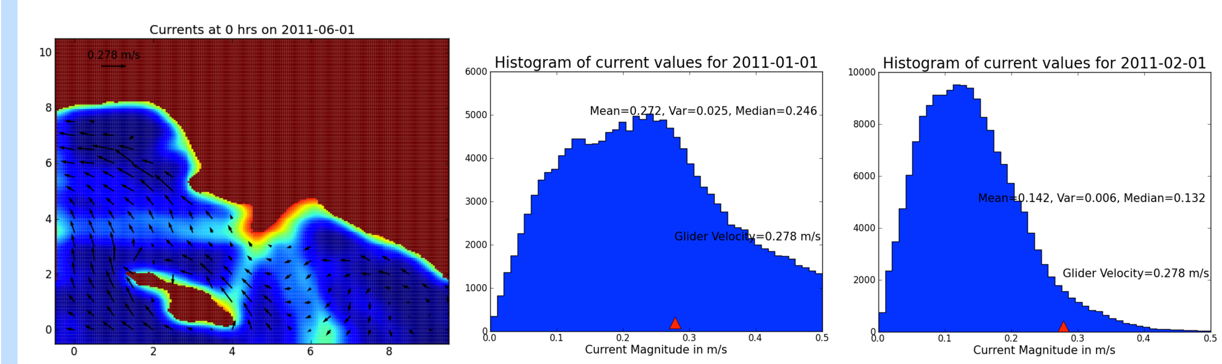

MDP vs. MER Replanner: Simulation and Experiments

- We developed Markov Decision Process (MDP) planner and Min-Expected Risk (MER-RP) Re-Planner
- Planners compute policies which minimize expected risk of surfacing while achieving goal
- Compared performance of planners in simulation and real experiments in the ocean
- MER-RP is more *goal-directed*
- MDPs provide general framework for *balancing goal-directed vs. risk-averse behavior*

G. Hollinger, A. M. Pereira, V. Ortenzi, and G. S. Sukhatme, Towards improved prediction of ocean processes using statistical machine learning, *Proc. Robotics: Science and Systems Workshop on Robotics for Environmental Monitoring (RSS)*, Sydney, Australia, July 2012.

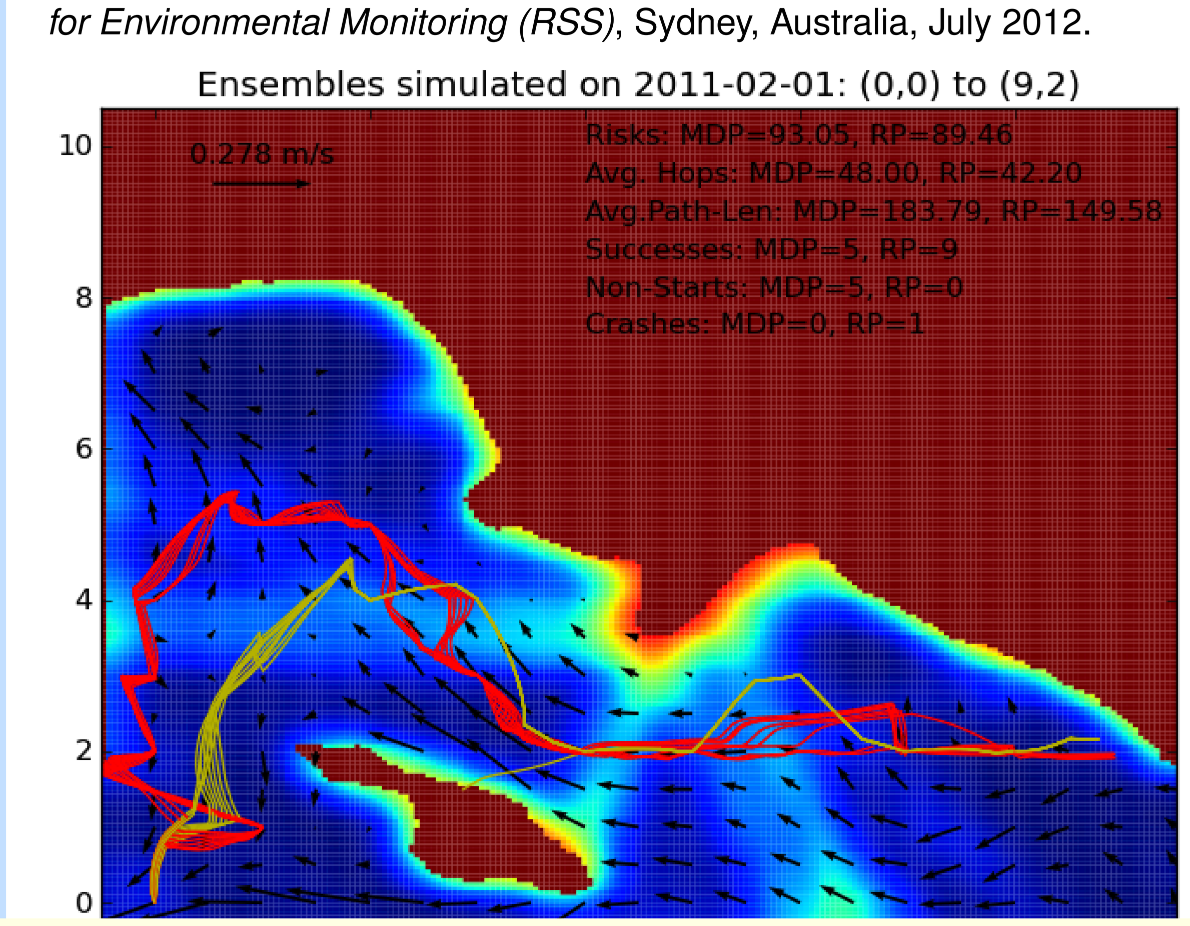

This work is supported by NSF CPS Medium Project 1035917, start date 10/10 through 9/13.