ONR Maritime Sensing - Discovery & Invention (D&I) Review
Naval Surface Warfare Center, Carderock (MD)
August 22-24, 2017

Multi-scale Adaptive Sensor Systems

Silvia Ferrari
Mechanical and Aerospace Engineering
Cornell University

Introduction

Distributed Optimal Control: the macroscopic state of the agents is represented by a restriction operator, such as a probability density function (PDF), to determine the optimal control laws for multi-agent systems (past work under Code 321).

- Multi-agent systems: few to hundreds of systems; heterogeneous; advanced sensing and, possibly, communication capabilities.
- Distributed control laws: path planning; obstacle avoidance; must meet one or more common goals, subject to agent constraints and dynamics.
- Derived and demonstrated DOC optimality conditions and algorithms.

Multi-scale Adaptive Control: a system of multiple autonomous dynamic systems that communicate and interact must adapt at different scales to cope with environmental changes and achieve evolving mission goals (new work under Code 321).

- Adaptation: manage control multiple assets and resources in the presence of significant uncertainties that cannot be modeled a priori.
- Multi-scale information gathering: individual assets can typically obtain high-quality in-situ measurements such that information can be fed back through the sensor and used to explain performance degradation.

Background: Distributed Optimal Control

- Agents' operating in Region of Interest (ROI) $W \subset \mathbb{R}^2$
- Performance measured in terms of restriction operator $\wp(\mathbf{x}_i,t)$
- Restriction Operator $\wp: W \times \mathbb{R} \rightarrow \mathbb{R}$
  - Time varying PDF $\wp(\mathbf{x}_i,t)$
  - PDF-based control law $\mathbf{u}_i\,(t)$

Terminal Cost

Instantaneous cost (Lagrangian)

$$J = \overset{\downarrow}{\phi[\wp(\mathbf{x}_i, T_f)]} + \int_{T_0}^{T_f}\int_X \overset{\downarrow}{\mathscr{L}[\wp(\mathbf{x}_i,t),\mathbf{u}_i(t),t]}\,d\mathbf{x}_i\,dt$$

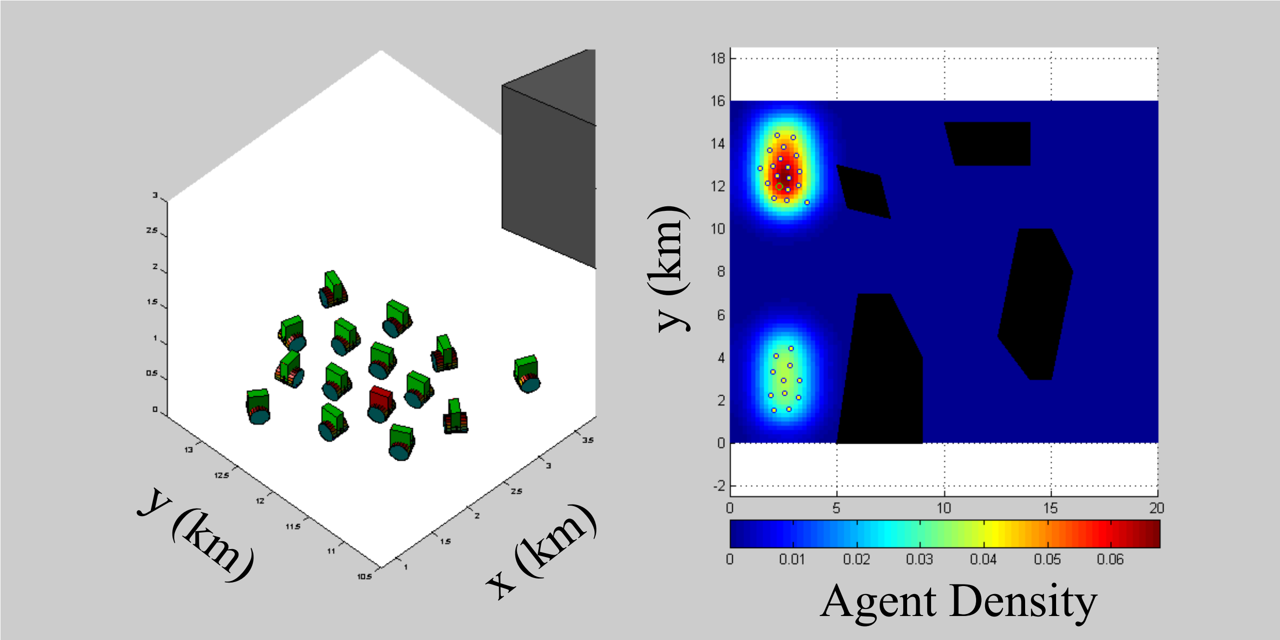

Background: Decentralized Sensing

Year 1: Communication Control for Active Sensing

Modeling of a spatial phenomenon, $g(\mathbf{x})$, by four robots with disjoint workspaces:

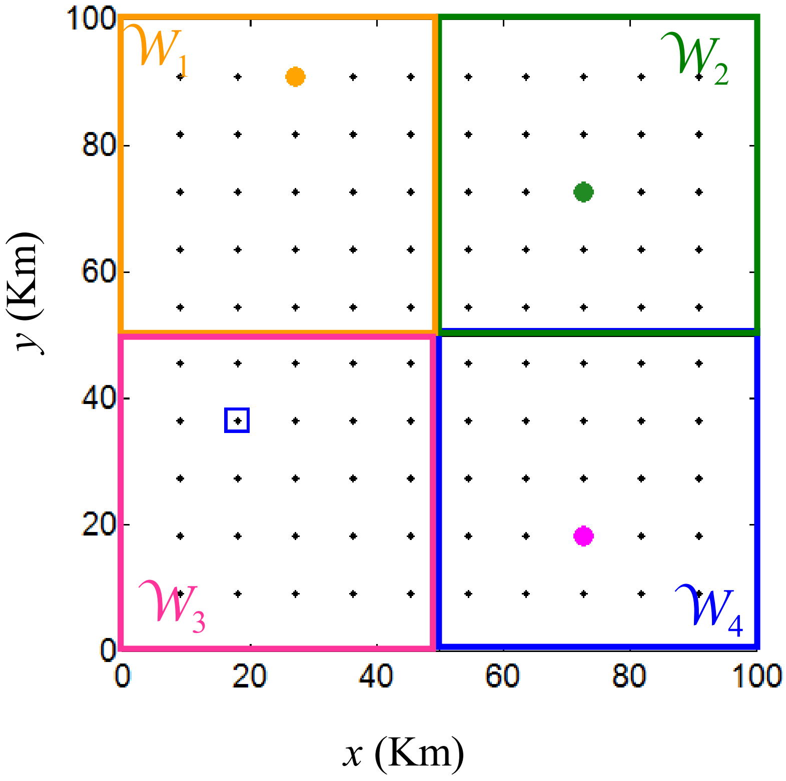

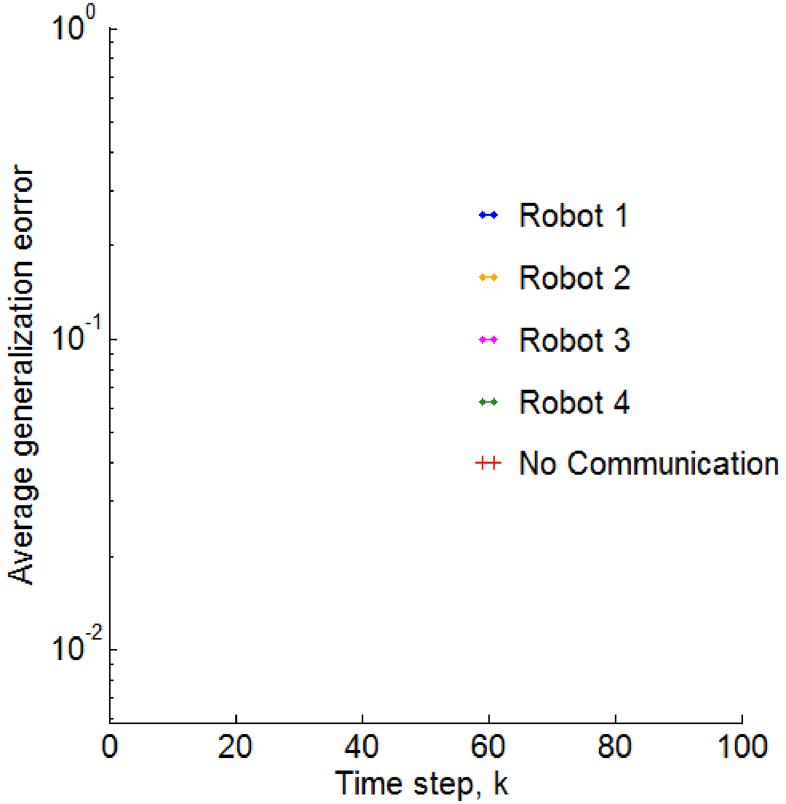

Motivation: Multi-scale Adaptive Control

- Environmental and operating conditions *in situ* may drastically differ those used *a priori*
  - – Off-line DOC solutions no longer optimal
  - – Agents must react to local information, including new tactical constraints
  - – Network-level controller can dispatch agents to localize gradient intensification while providing energy management, volume coverage, and robustness to component failure

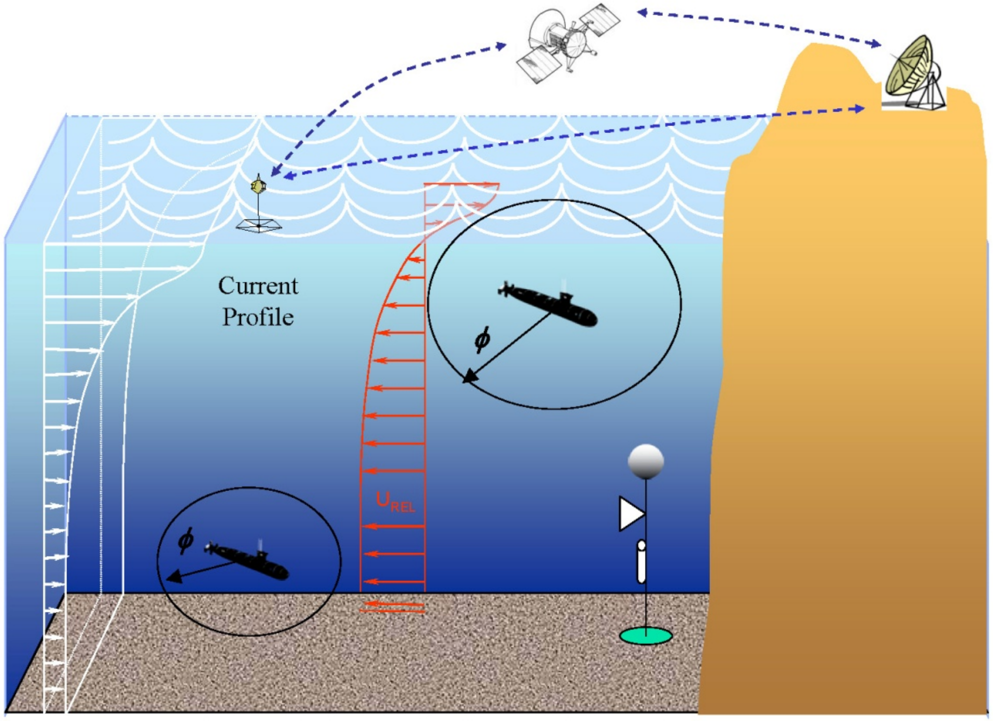

1. **Target:** actual population is different from that assumed *a priori*.
2. **Environment:** conditions measured in situ are different from those forecasted by oceanographic models.
3. **Platform:** navigation settings are suboptimal, leading to incorrect estimates of agent position and/or direction.
4. **Sensor:** actual performance is different from the performance function model due to the above conditions, or sensor malfunctioning.

Motivation: Multi-scale Adaptive Control

- Environmental and operating conditions *in situ* may drastically differ those used *a priori*
  - Off-line DOC solutions no longer optimal
  - Agents must react to local information, including new tactical constraints
  - Network-level controller can dispatch agents to localize gradient intensification while providing energy management, volume coverage, and robustness to component failure

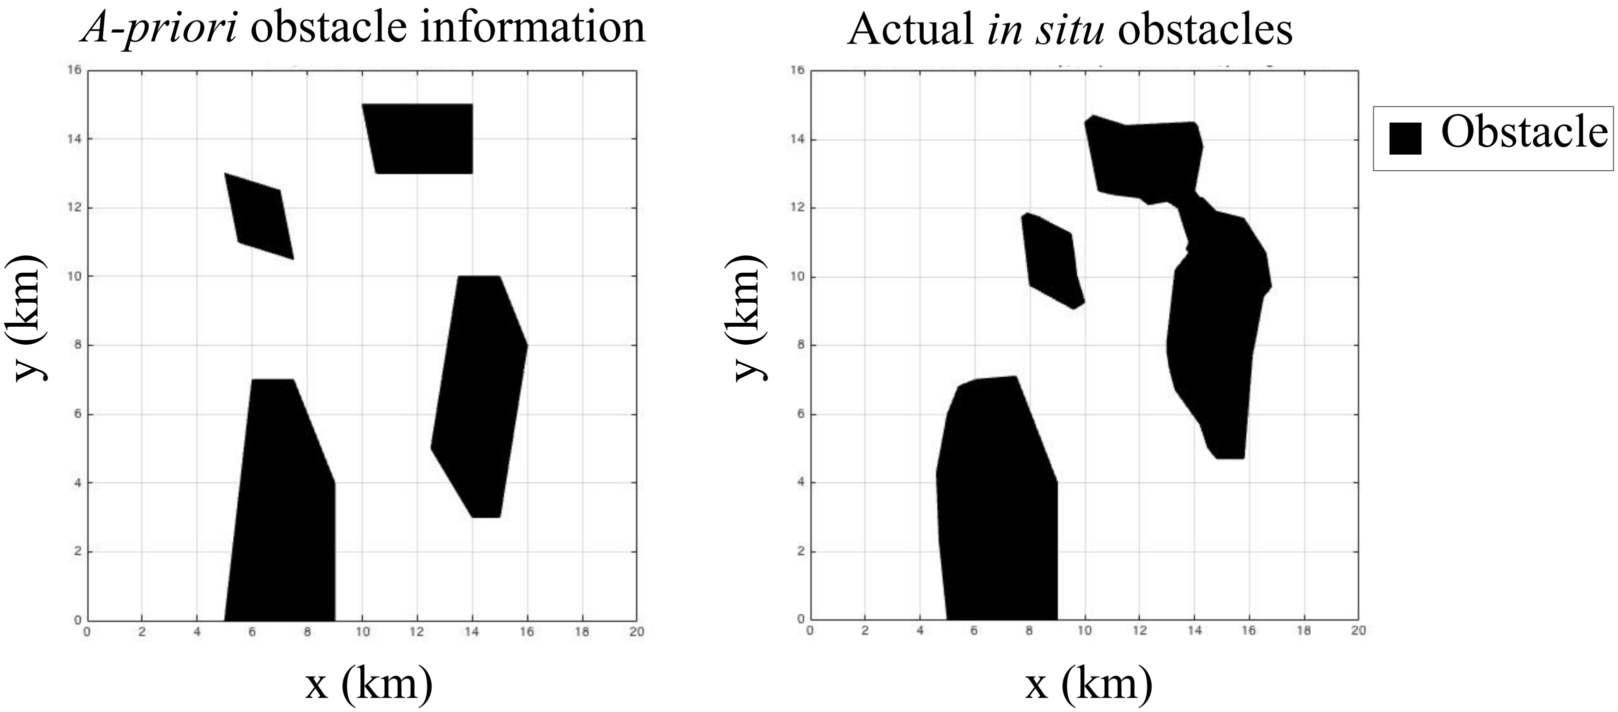

Motivation: Oceanographic Conditions

- Full, nonlinear dynamics $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}, t)$
- Robot state: $\mathbf{x} = [\underbrace{x, y, z}_{\text{Inertial position}}, \underbrace{\theta, \psi}_{\text{Pitch and yaw Euler angles}}, \dot{x}, \dot{y}, \dot{z}, \dot{\theta}, \dot{\psi}]^T$

•Robot Control: $\mathbf{u} = [\delta_{rpm}, \delta_r, \delta_s]$

- UUV kinematics:

$$\dot{x} = \|\mathbf{v}_B\|(\cos\theta\cos\psi) + v_{cx}$$
$$\dot{y} = \|\mathbf{v}_B\|(\cos\theta\sin\psi) + v_{cy}$$
$$\dot{z} = \|\mathbf{v}_B\|(-\sin\theta)$$
$$\dot{\theta} = g_\theta \delta_r$$
$$\dot{\psi} = g_\psi \delta_s$$

$\mathbf{v}_B$: measured velocity
v_{cx}, v_{cy}: measured ocear
and y velocities
g_θ, g_ψ : control gains

A-priori current estimates:

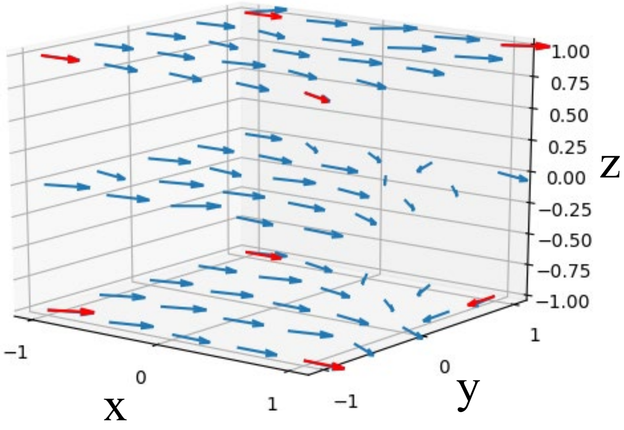

Actual currents:

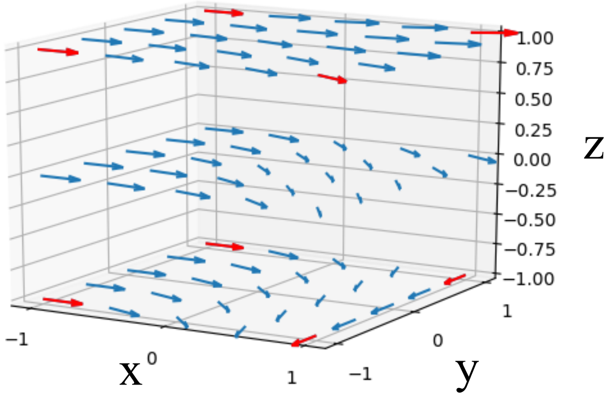

Problem Formulation

Problem Formulation

Mission Goal: Collectively *explore* and *map* obstacles and currents in a region of interest *W* while obtaining decentralized sensor measurements, avoiding obstacles, and communicating with other agents and a central station.

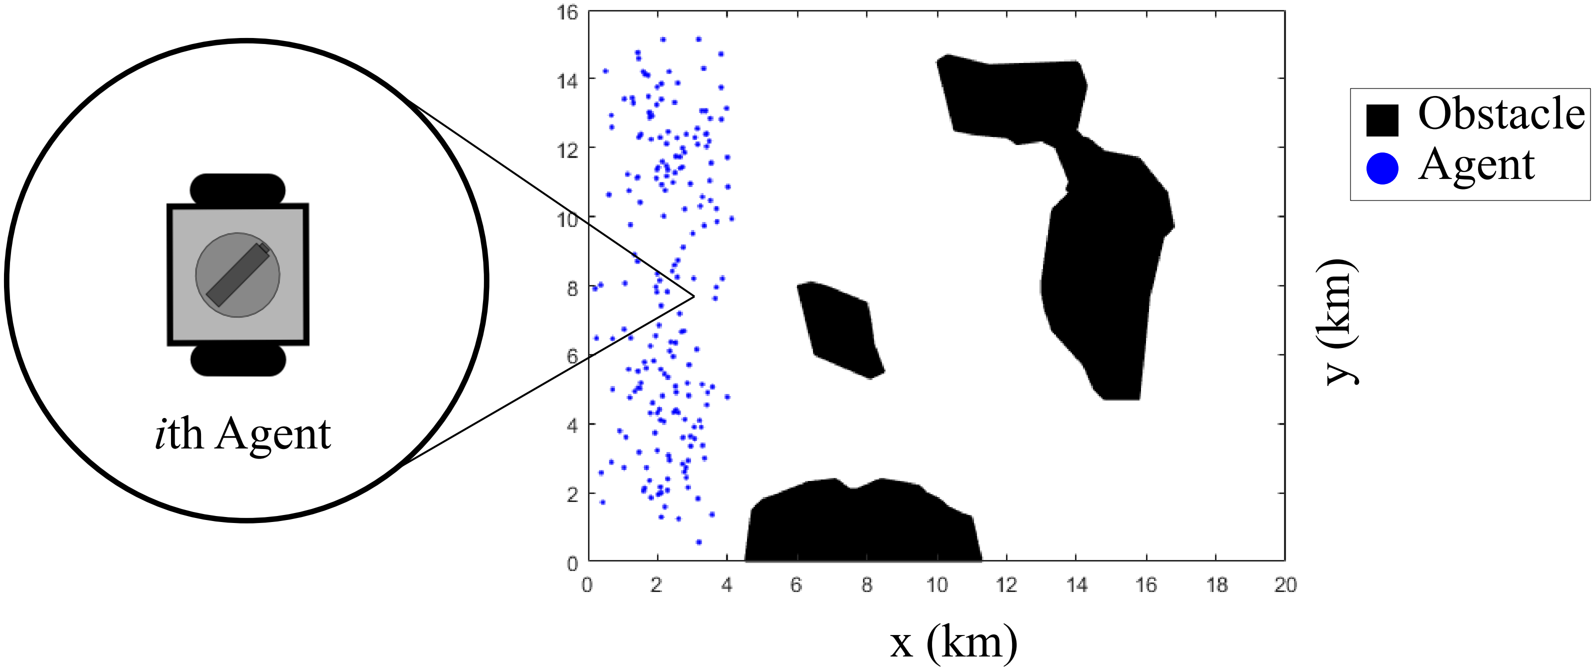

- Region of Interest $W \subset R^2$:
- Fixed, unknown, rigid obstacles, B_i , $i=1,\ldots,r$
- N agents

On-board Sensor Measurements

- Sensor can infer classification $\hat{Y}$ within sensor range and construct $D_i = \{\mathbf{x}_j, y_j\}^m$

- Noisy sensor measurements:

$$\mathbf{x}_M = \mathbf{x}_i + D\hat{\mathbf{e}}_r$$

$$\hat{Z} = [\hat{\mathbf{x}}_M, \hat{Y}]$$

- Y is a random and binary classification variable:

$$Y(\mathbf{x}) = \begin{cases} 1, \text{if } \mathbf{x} \in B \\ 0, \text{if } \mathbf{x} \notin B \end{cases}$$

$\mathrm{D} = d + v$: Distance measurement
$\Theta = \theta + \upsilon$: Angle measurement
$v \sim N(0,\sigma_d)$: Sample from normal distribution
$\upsilon \sim N(0, \sigma_\theta)$: Sample from normal distribution
$\hat{\mathbf{e}}_r = \hat{x}\cos\Theta + \hat{y}\sin\Theta$: radial unit vector
$\hat{x}, \hat{y}$: unit vectors of basis in F_A

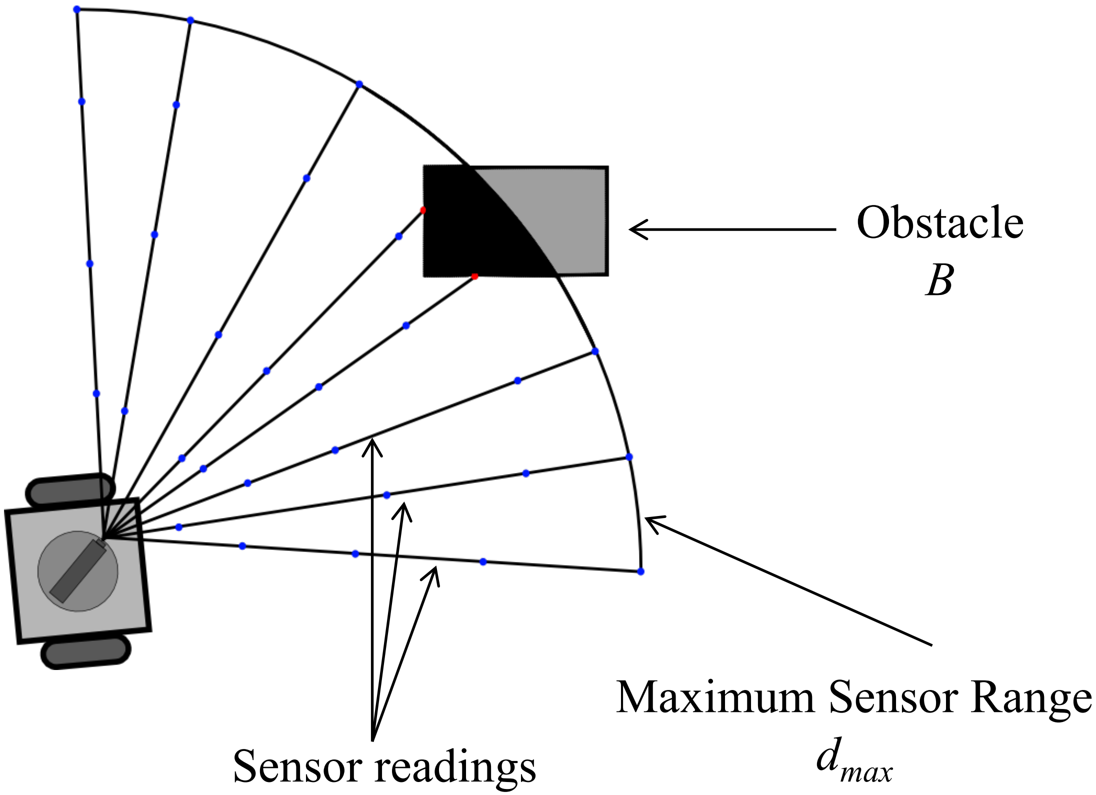

🟥 Hit/Occupied
🟦 Miss/Not Occupied

On-board Sensing and Communications

- **Sensing goal:** maximize information gain of future measurements to minimize uncertainty
- **Mission constraints:**

  Bounded sensor FOV range, R

  Bounded communication range, R_c
- **Multiobjective Cost Function:**

$$J = \int_{W} \hat{\varphi}(Y;Z\,|\,M,\lambda)d\mathbf{x} + \int_{T_0}^{T_f} \left\{U[\mathbf{x}_i(t)] + \mathbf{u}_i(t)^T\,\mathbf{R}\mathbf{u}_i(t)\right\}dt$$

Bayesian measurement model: $p(Z\,|\,Y,\lambda)$

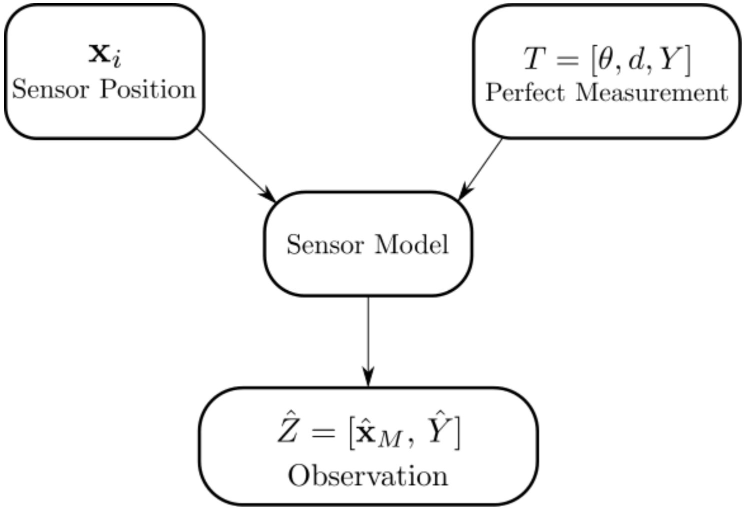

$\mathbf{x}_i$ Sensor Position

$T = [\theta, d, Y]$ Perfect Measurement

Sensor Model

$\hat{Z} = [\hat{\mathbf{x}}_M, \hat{Y}]$ Observation

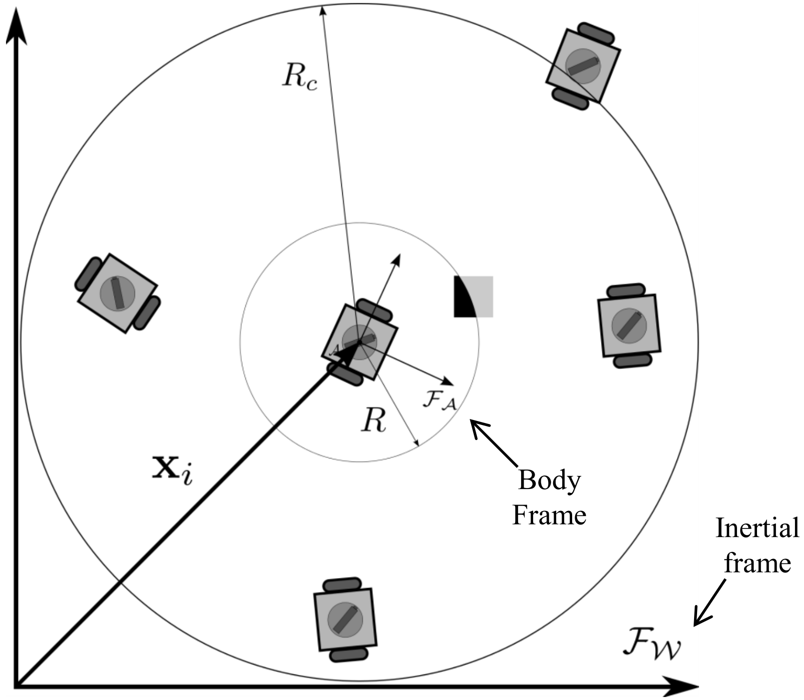

- $\hat{\varphi}(\cdot)$: Information gain
- Y, Z: hidden discrete random variables
- M: set of all prior measurements
- λ: environmental condition parameters
- $U(\mathbf{x}_i)$: obstacle repulsion potential
- R: control weight matrix

Technical Approach

Hilbert Mapping

- Advantages of **Hilbert mapping** for multi-scale systems:
  - Continuous
  - Probabilistic
  - Spatial correlations preserved
- Nonlinear mapping problem ~ *binary classification task*
- Approximate probability of occupancy as,

$$P(Y=1 \mid X=\mathbf{x};\mathbf{w}_i)=1-\frac{1}{e^{\mathbf{w}_i{}^T\Phi(\mathbf{x})}}$$

by learning vector of parameters $\mathbf{w}_i$ online.

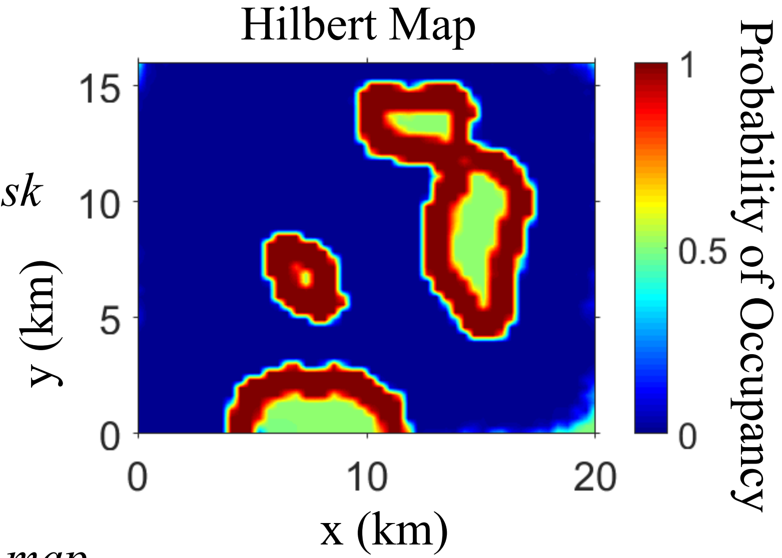

- $\Phi(\mathbf{x})$: $\mathbb{R}^2 \rightarrow \mathbb{R}^n$ known as a *lifting* function or *feature map*
- From Mercer's theorem, for any non-negative definite function, K(**x**, **x**'), there exists $\Phi(\mathbf{x})$ such that: $K(\mathbf{x},\mathbf{x}')=\langle\Phi(\mathbf{x}),\Phi(\mathbf{x}')\rangle$
  - This is known as a *kernel function*
  - **Example:** Radial Basis Function Kernel: $K(\mathbf{x},\mathbf{x}')=\exp(-\frac{1}{2\sigma^2}\left\|\mathbf{x}-\mathbf{x}'\right\|^2)$

“Kernel Trick” in Hilbert Mapping

- Define parameters as: $\mathbf{w}_i = \sum_{j=1}^{\mu_i} \alpha_j \Phi(\mathbf{x}_j)$, where μ_i is the dimension of the data set
- Therefore, the function to be learned from sensor data is:

$$P(Y=1 \mid X=\mathbf{x}_*;\mathbf{w}_i) = 1 - \frac{1}{e^{\mathbf{w}_i{}^T \Phi(\mathbf{x}_*)}} = 1 - \frac{1}{1+\exp\left(\sum_j^{\mu_i} \alpha_j \Phi(\mathbf{x}_j)\cdot\Phi(\mathbf{x}_*)\right)}$$

$$= 1 - \frac{1}{1+\exp\left(\sum_j^{\mu_i} \alpha_j K(\mathbf{x}_j,\mathbf{x}_*)\right)}$$

“Kernel Trick”

- Parameters α$_i$ determined by minimizing a regularized loss function

$$\min_{\vec{\alpha}} L(\vec{y}, K\vec{\alpha}) + \lambda \vec{\alpha}^T K \vec{\alpha}$$

where $\vec{y}$ is a vector composed of y_j from data set *D, K* is a kernel matrix, formed by applying the chosen kernel function to *D,* and λ is a regularization parameter (scalar).

- Negative Log Likelihood Loss function ~ find α$_j$ by maximum likelihood estimation (MLE)

Example: Kernel Methods

Data not linearly separable in $\mathbb{R}^2$

$\mathbf{x}$ =[x, y]

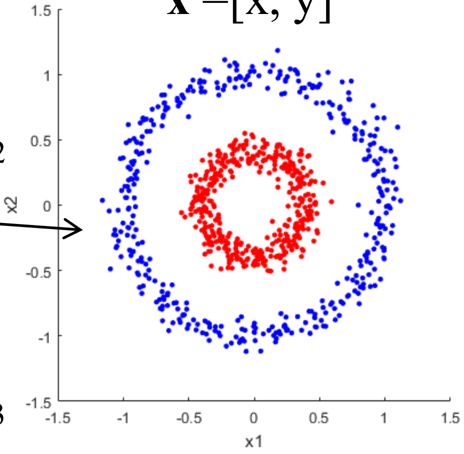

$\Phi(\mathbf{x})$

$[x, y] \rightarrow [x, y, x^2 + y^2]$

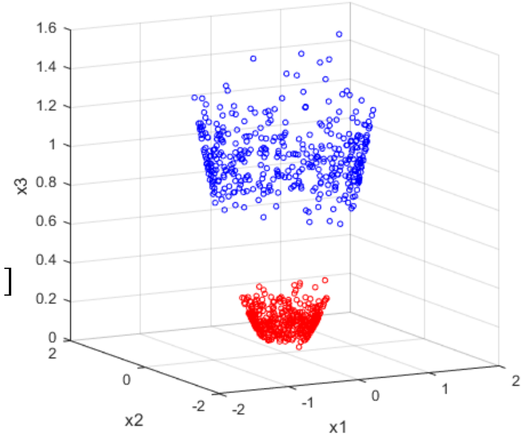

Nonlinear decision boundary

Linearly separable in $\mathbb{R}^3$

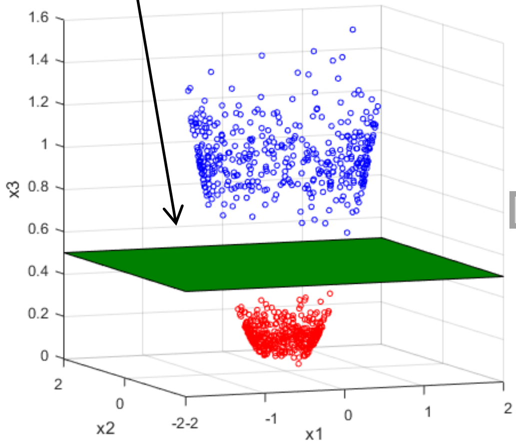

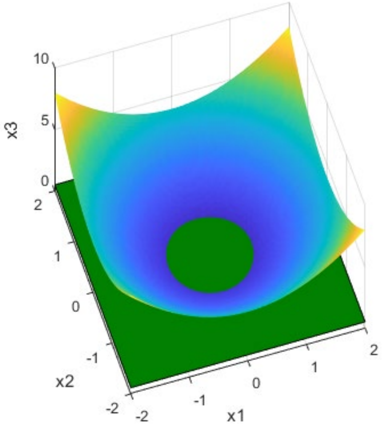

Project back to original dimension

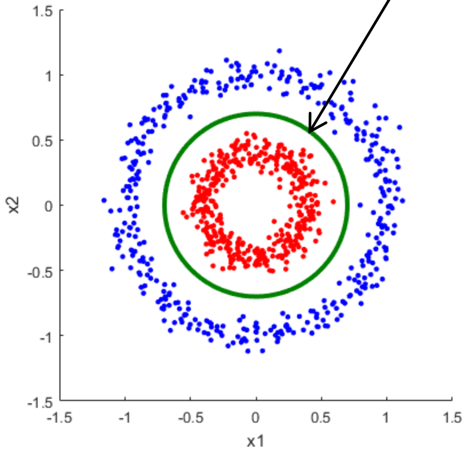

Local Hilbert Mapping

- Agent i stores data set $D_i = \{\mathbf{x}_j, y_j\}^m$, $j=1,\dots,m$, obtained while navigating in ROI
  - $\mathbf{x}_j \in R^2$
  - $y_j \in \{0,1\}$ categorical variable

Sensor Data Collected

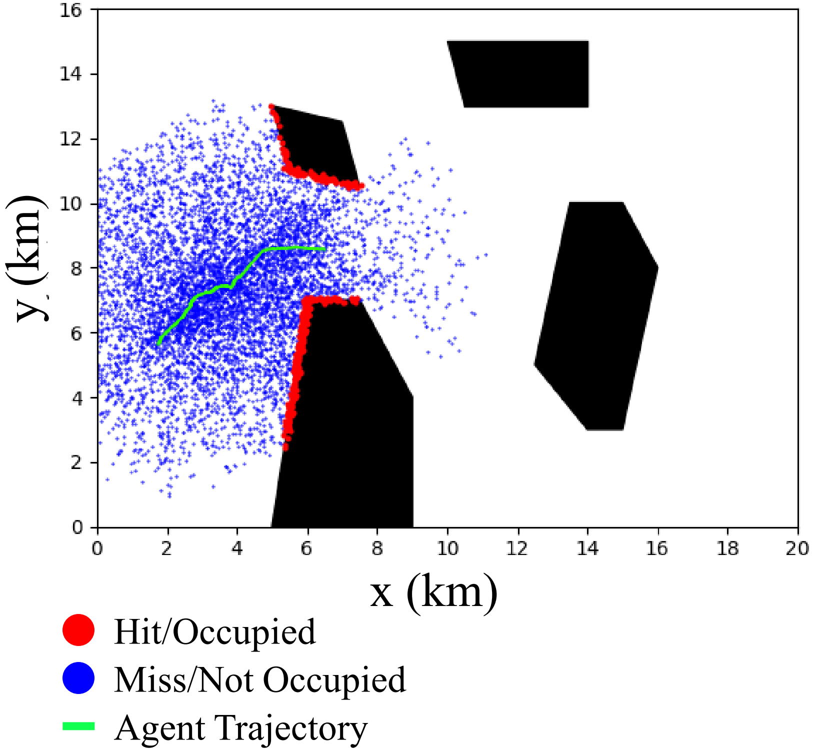

Resulting Hilbert Map

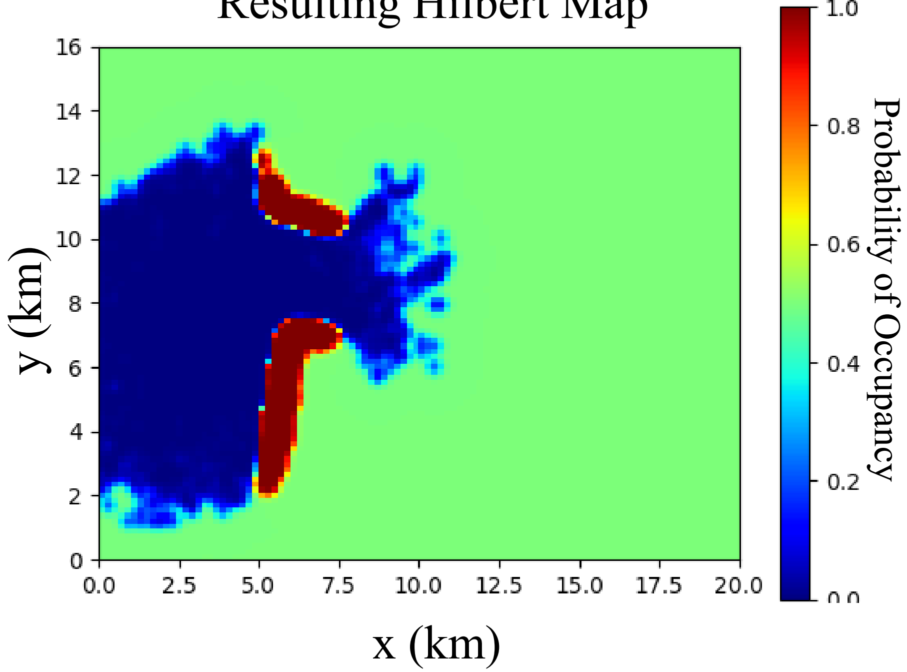

Decentralized Hilbert Map

Global Hilbert Mapping:

- Compute Gaussian Mixed Model (GMM) centers from sensor data:

$$\tilde{p}(t) = \sum_{k=1}^{M} w_k \mathcal{N}(\mu_k \,|\, \Sigma_k)$$

where M is the number of components, w_k is a scalar mixing coefficient, and $\mathcal{N}(\mu, \Sigma)$ is a bivariate normal distribution with mean μ, and covariance matrix Σ.

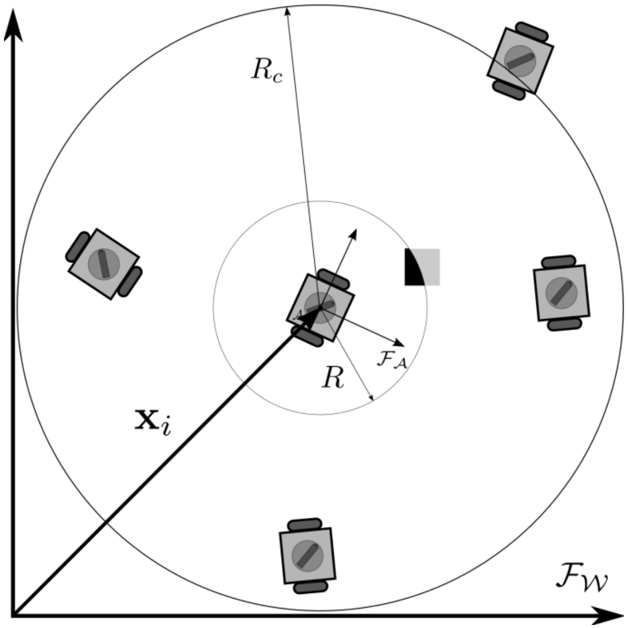

For each agent, labeled by $i = 1, \ldots, N$:

- Communicate $S=\{\mu_k\}_{k=1,\ldots,M}$ to neighbors in R_c
  - – Communicating S in lieu of D requires *less* memory
  - – Kernel Methods have shown to train classifiers using less training data than the massive amount of data originally collected by sensors
- Other agents incorporate S into own training data set and update their maps

Decentralized Information Sharing

Assumptions:

- Connected agents share GMM centers
- Agents share information with all those connected to their network (comm protocol)

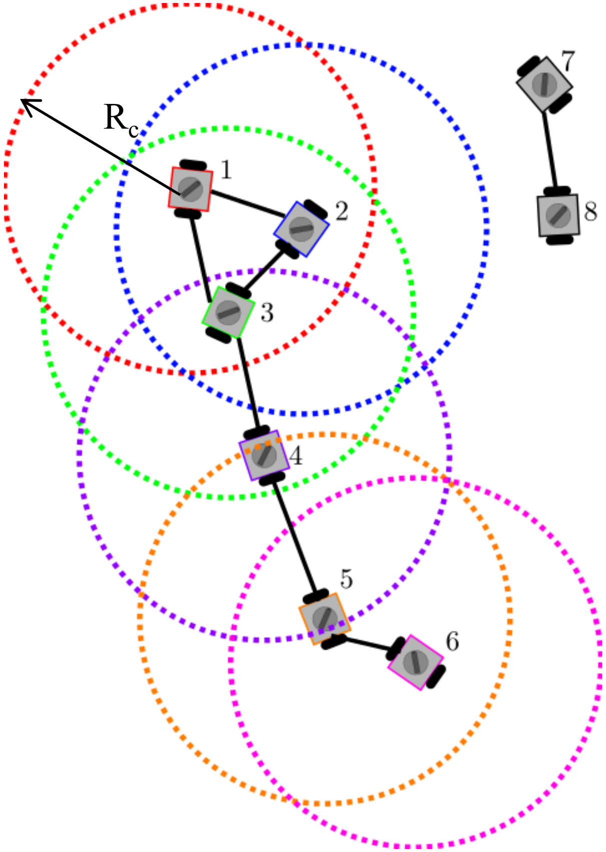

Hilbert Map:

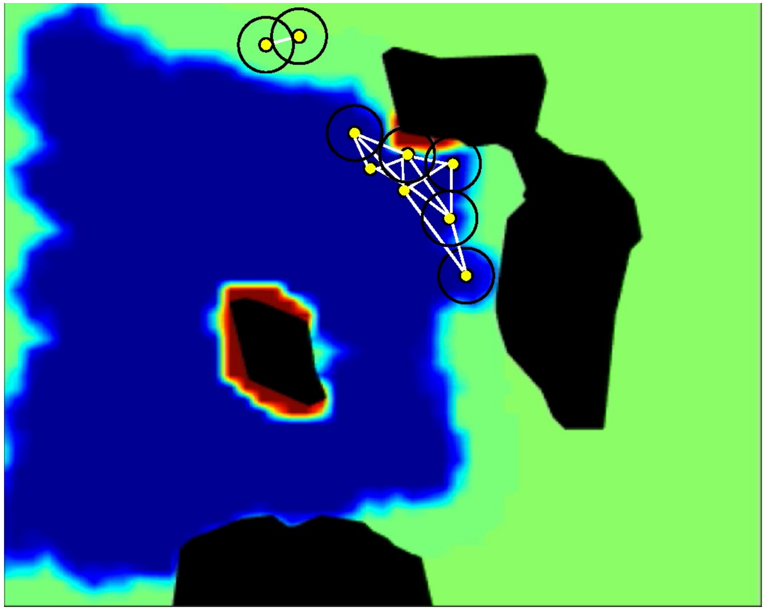

Communication Graph

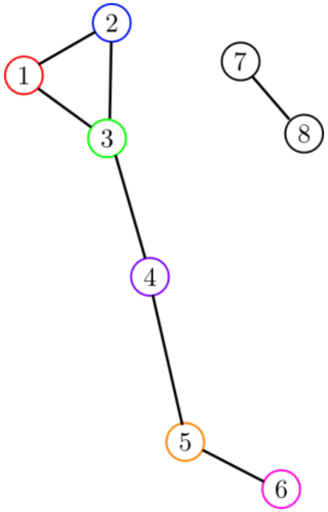

Decentralized Hilbert Map Results

- Agents gain information in regions far outside their FOV
- Agents rely on local map when communications are unavailable

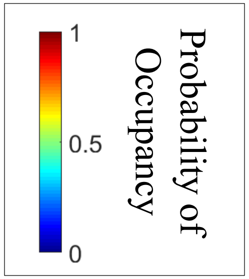

Agents' Maps at $t = 4$, before communication

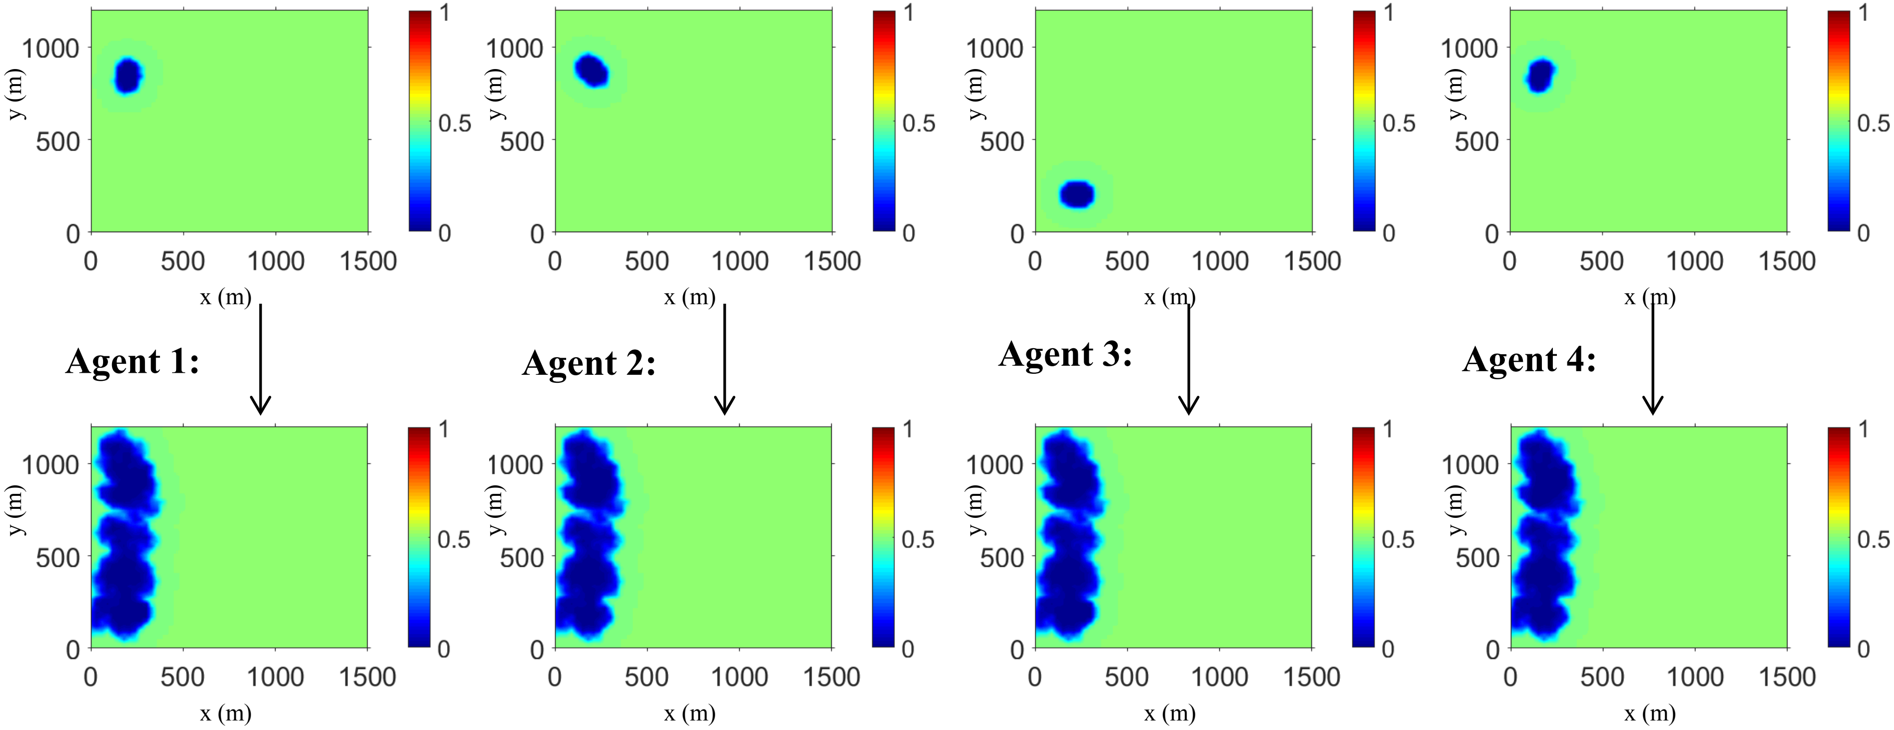

Agents' Maps at $t = 5$, after communication with all *connected* agents

Hilbert Map Information Value

- ***Information value*** of future sensor measurements is used here for planning
- Expected entropy reduction: reduction in uncertainty caused by measurement Z_k

$$\Delta H(Z_k) = \underbrace{H(Y \mid Z_{k-1}, \lambda)}_{\text{Conditional Entropy}} - \underbrace{E(H \mid Z_k, \lambda)}_{\text{Expected Entropy after measurement } Z_k}$$

- All terms obtained from Bayesian model:

$$P(Y \mid M_{k-1}, \lambda) = \frac{\overbrace{P(Z_{k-1} \mid Y, \lambda)}^{\text{Measurement model}}\overbrace{P(Y \mid M_{k-2}, \lambda)}^{\text{Hilbert Map}}}{\sum_{y \in \mathscr{Y}} P(Z_{k-1} \mid Y = y, \lambda) P(Y = y \mid M_{k-2}, \lambda)}$$

$$P(Y \mid Z = z, M_{k-1}, \lambda) = \frac{\overbrace{P(z \mid Y, \lambda)}^{\text{Measurement model}}\overbrace{P(Y \mid M_{k-1}, \lambda)}^{\text{Hilbert Map}}}{\sum_{y \in \mathscr{Y}} P(z \mid Y = y, \lambda) P(Y = y \mid M_{k-1}, \lambda)}$$

- $H(\cdot)$: Entropy
- Y, Z: discrete random variables representing occupancy of an area and measurements at time k
- $M = \{Z_1 \ldots Z_k\}$: set of all previous measurements
- λ: environmental condition parameters
- $Y, Z_k \in \{0, 1\}$

Agent Path Planning

- Information Roadmap (IRM) Method [Zhang, Ferrari, ‘09]
  - Sample locations in W from cost function and uniformly around the agent
  - Connect nodes that go through “safe” areas based on Hilbert Map
  - Plan path over graph to nodes with highest information value

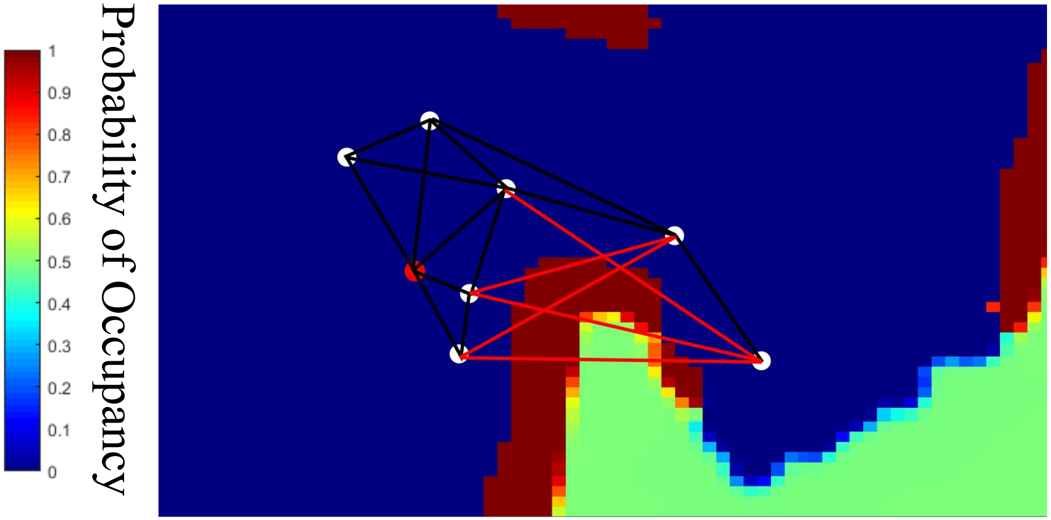

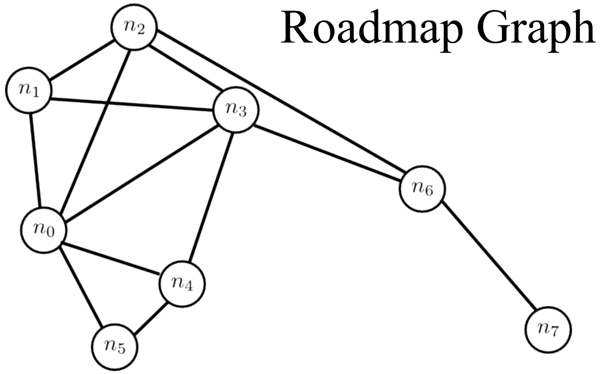

Information Roadmap Path Planning

- Nodes for the information roadmap are sampled from potential information function.
- Potential information function is created from map of information value (generated from Hilbert map) and obstacle repulsive potential.

Hilbert Map

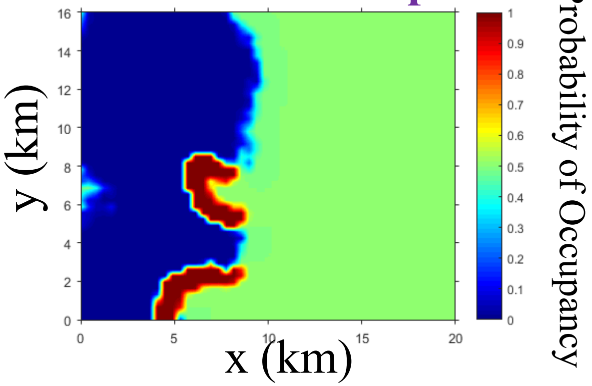

Information Value

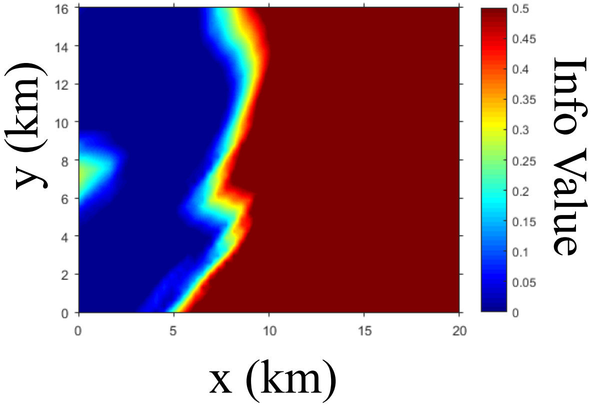

Black areas are *very large* penalties

Cost Function

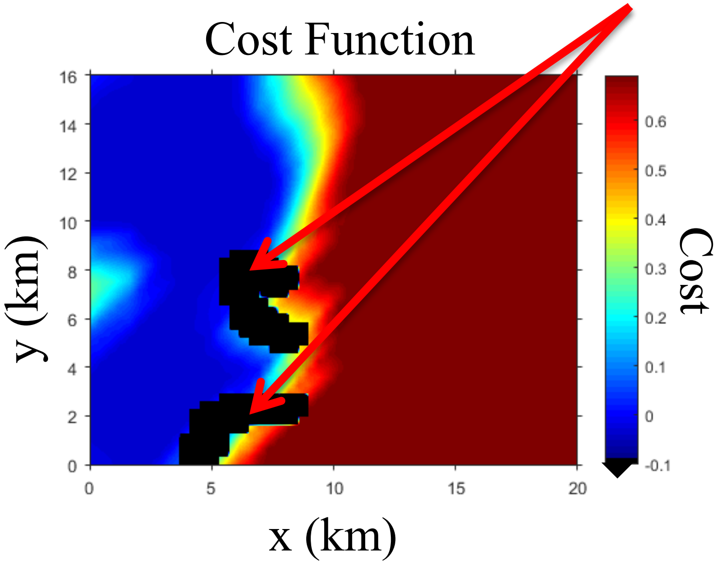

Simulation Results

Centralized Information-driven Planning

- Agents all share information and compute a centralized Hilbert Map
- Each agent plans own path using information roadmap method

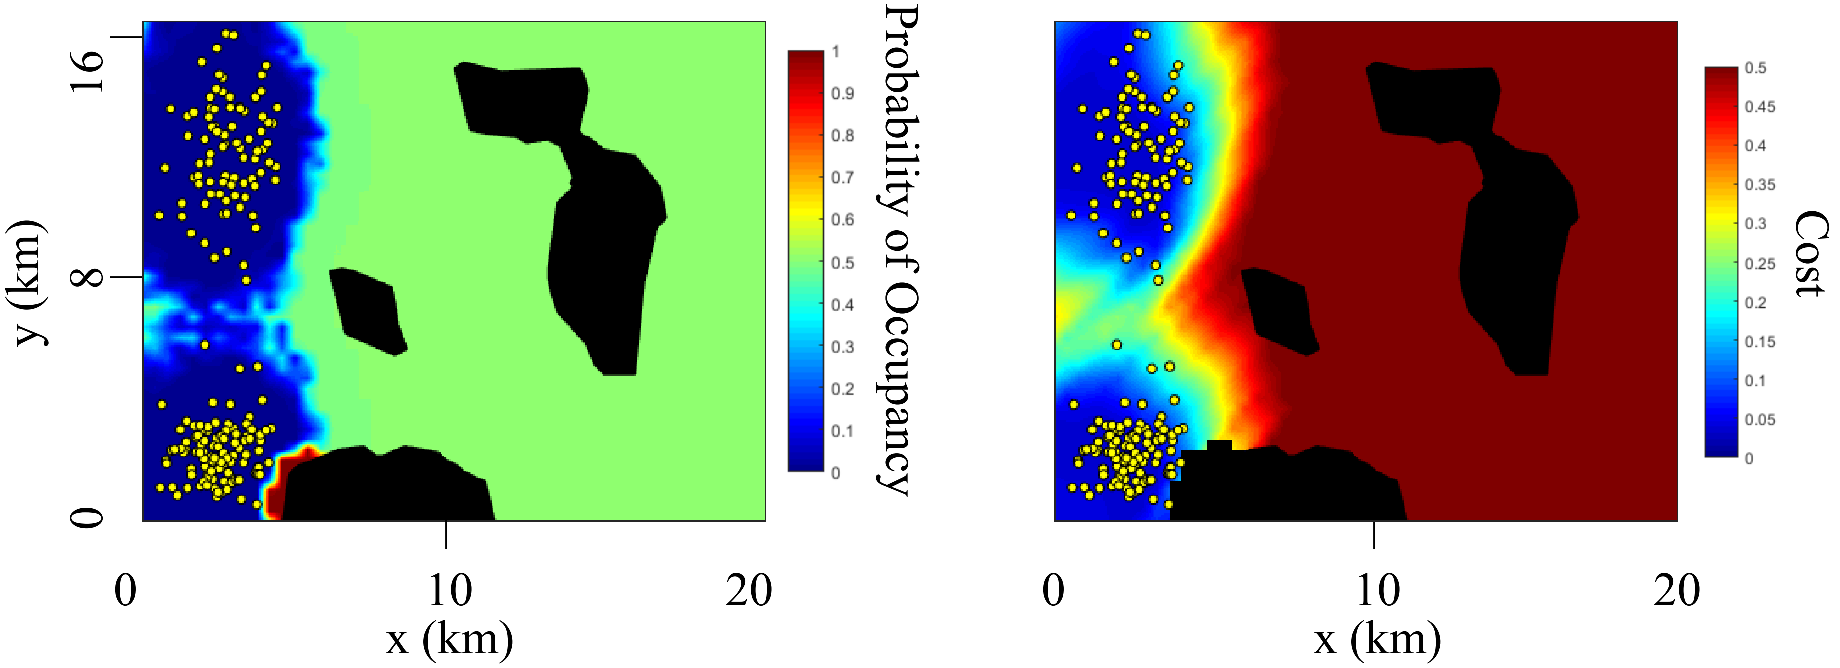

200 Agents
FOV = 1.5 km
Centralized fusion

Agent Information-driven Planning without Communications

- One agent exploring alone
- Unable to map large ROI in a reasonable amount of time

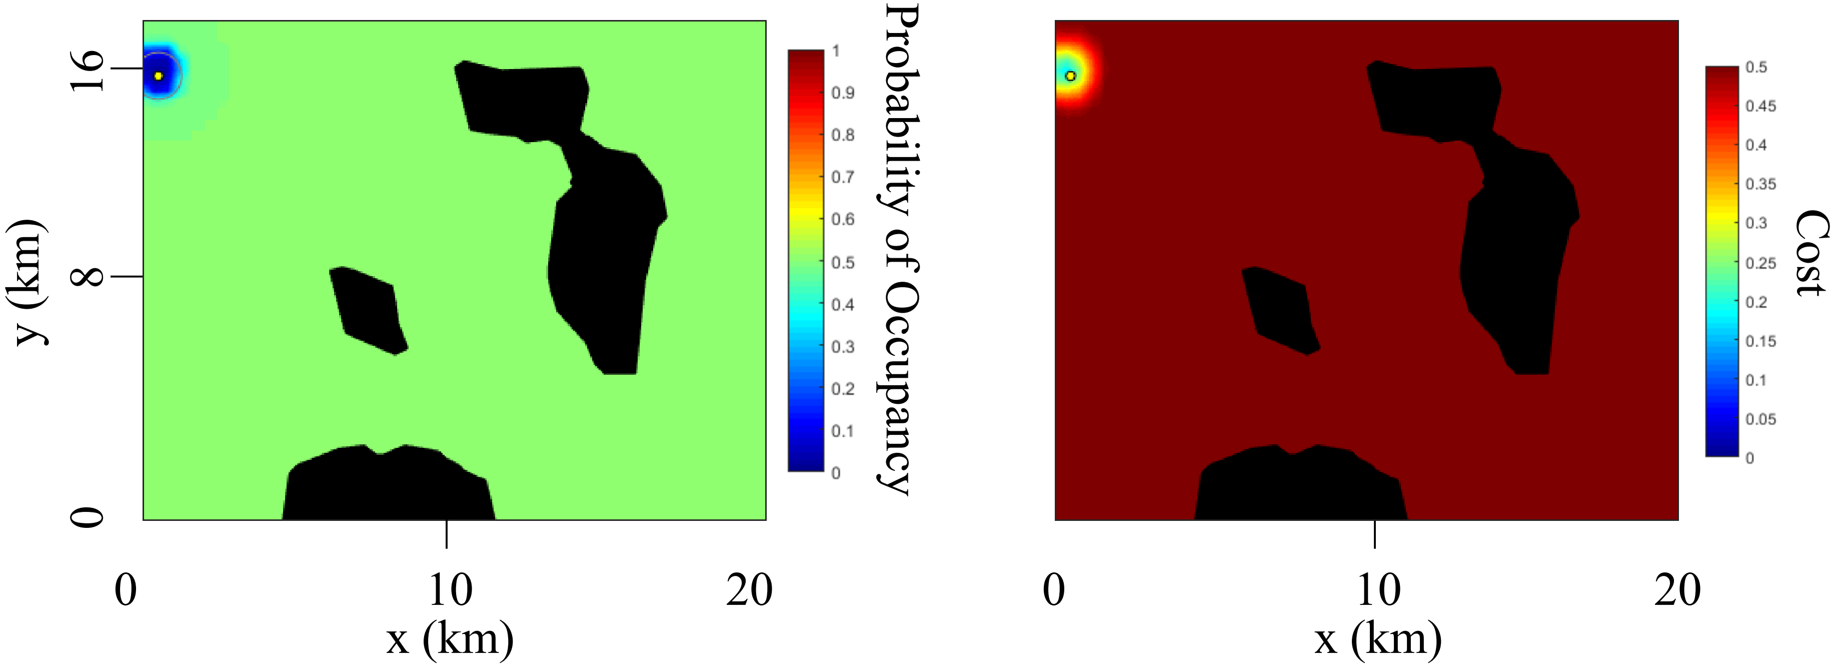

1 Agent
FOV = 0.75 km
No fusion

Agent Information-driven Planning with Communications

- Two sample agents in a network of communicating with neighbors in range (R_c)

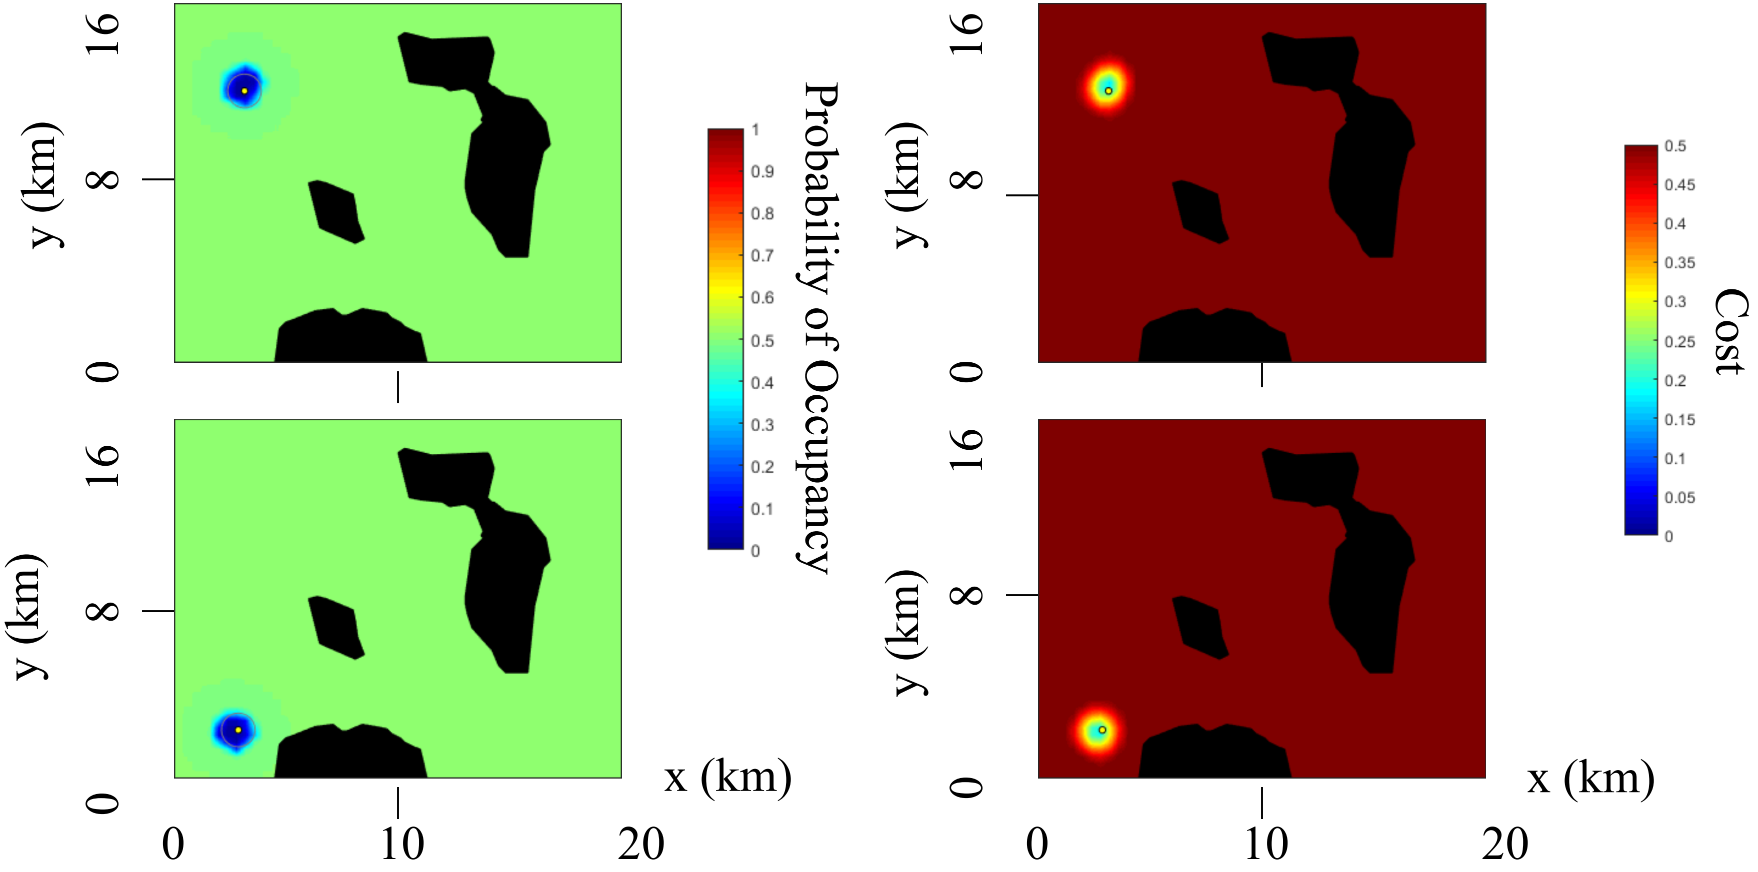

FOV = 0.75 km; N = 200 agents
Communication Range (R_c) = 4.5 km

Decentralized Communication Results

- Connected agents (yellow) achieve higher information value and avoid regions already explored by others.
- Disconnected (blue) agents “look lost” and converge to regions already mapped with high confidence.

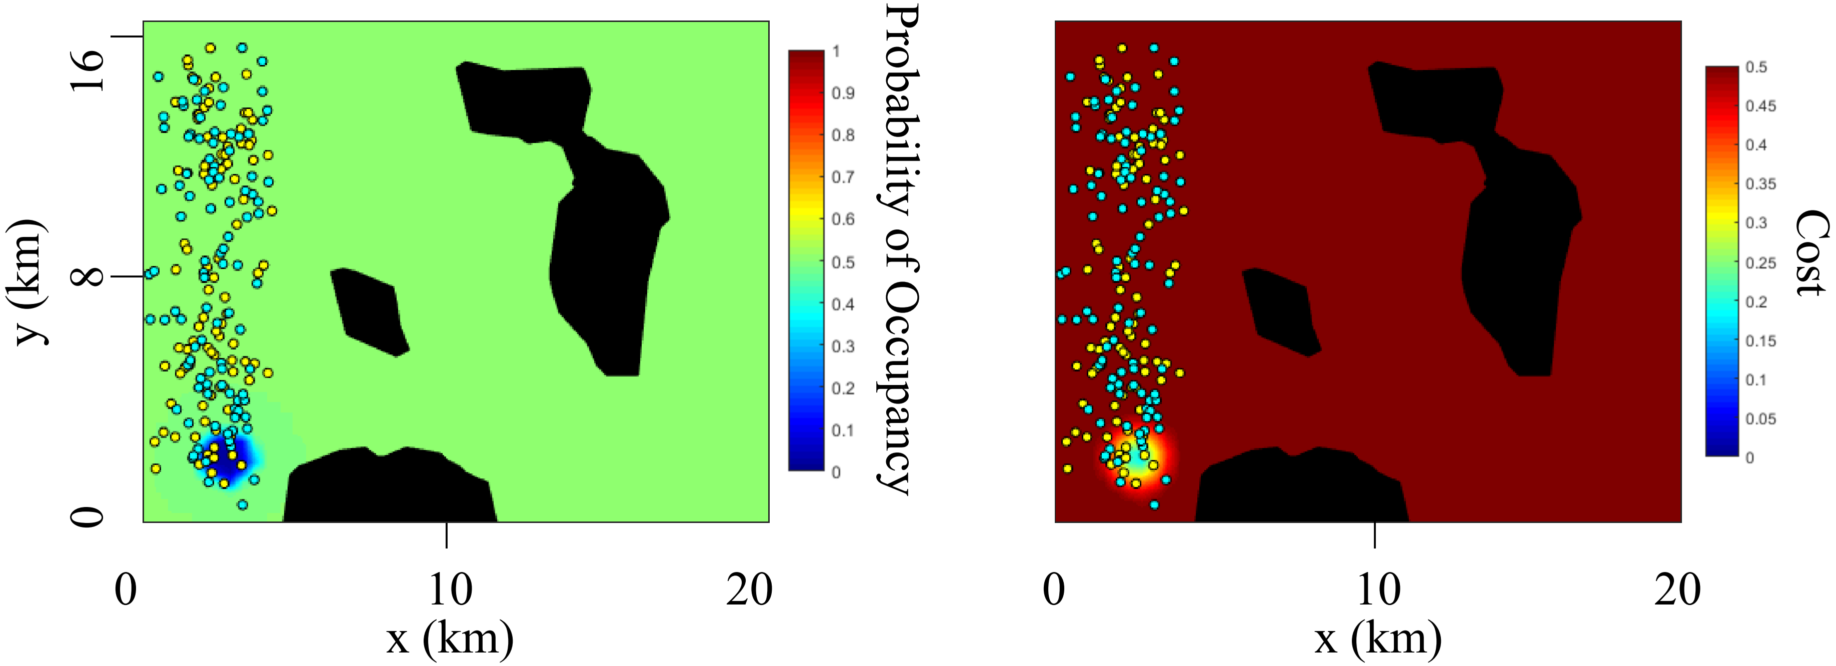

200 Agents, only yellow agents (50%) communicate
FOV = 0.75 km
Communication Range = 4.5 km

Local vs. Global Information Value

- A multiscale dynamical system with communications can far outperforms a single agent in distributed sensing tasks
- Next question: how does unmanaged information propagating through network affect system stability?
- Agent planning: respond both to the network objectives (based on global data set) vs. local information (unavailable to the network controller due to comm. delays).

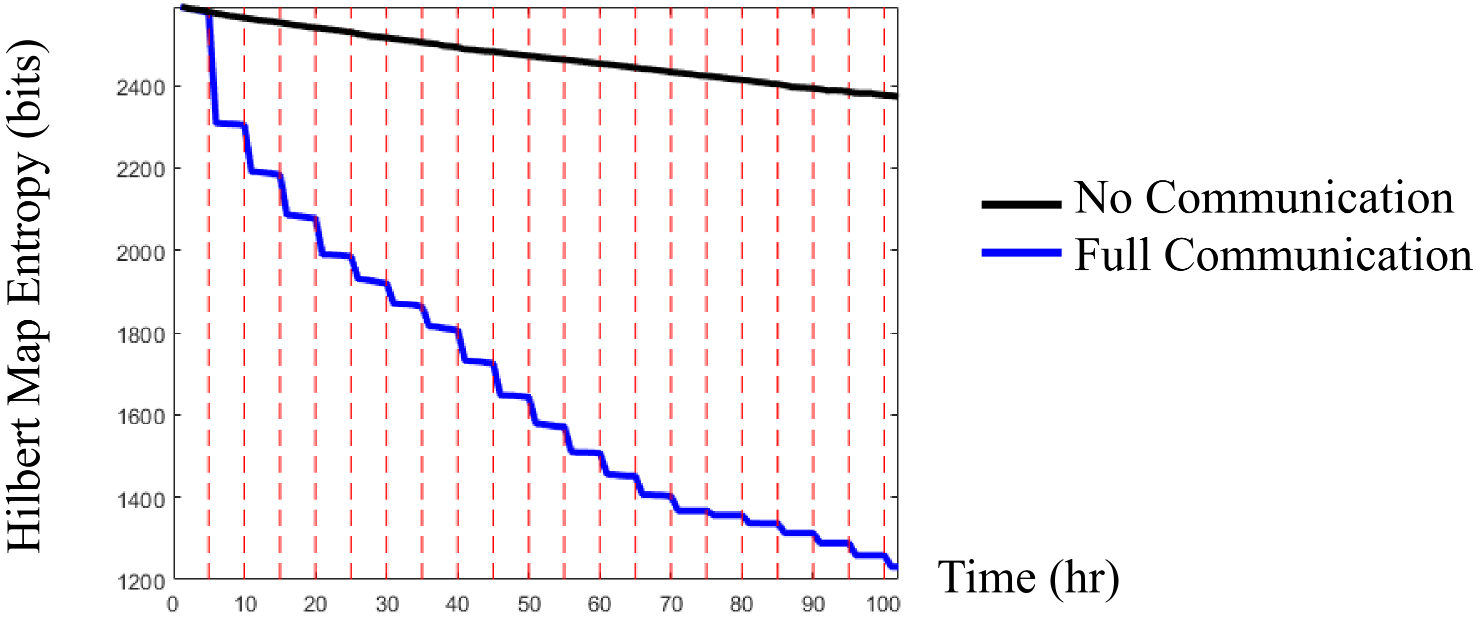

Multi-Scale Adaptive Optimal Control

- *Adaptive* distributed optimal control:
  - Agents make decisions based on *in situ* conditions and *global* information
  - Value function V, defined in terms of discrete $\wp_k$, at time k and control law $C(\wp_k)$:

$$V \equiv \phi[\wp_{k_f}] + \sum_{t_k}^{t_{f-1}} \int_X \mathscr{L}[\wp_k, C(\wp_k)] d\mathbf{x}_i = V[\wp_k, C(\cdot)]$$

- Control law C_l and Value function V_l are iteratively improved online, where l is the iteration
- $\wp_k$ agent density at discrete time k

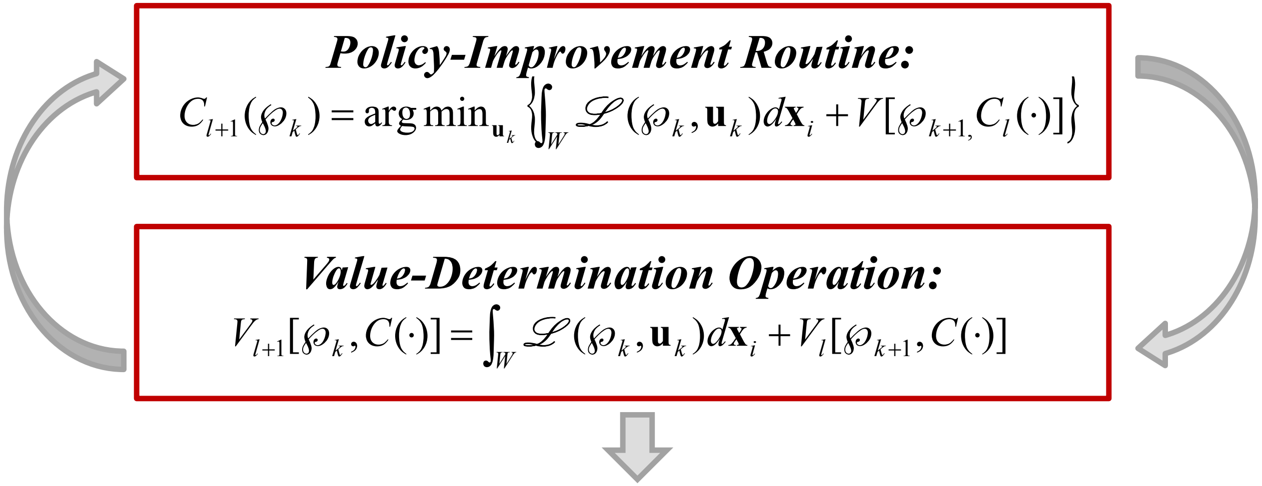

Policy-Improvement Routine:

$$C_{l+1}(\wp_k) = \arg\min_{\mathbf{u}_k} \left\{ \int_W \mathscr{L}(\wp_k, \mathbf{u}_k) d\mathbf{x}_i + V[\wp_{k+1,} C_l(\cdot)] \right\}$$

Value-Determination Operation:

$$V_{l+1}[\wp_k, C(\cdot)] = \int_W \mathscr{L}(\wp_k, \mathbf{u}_k) d\mathbf{x}_i + V_l[\wp_{k+1}, C(\cdot)]$$

Optimal value function V^*:

$$V^* = \min_{\mathbf{u}_k} \left\{ \int_X \mathscr{L}(\wp^*{}_k, \mathbf{u}_k) d\mathbf{x}_i + V^*[\wp^*{}_{k+1}, C^*(\cdot)] \right\}$$

Summary and Conclusions

Multi-scale Adaptive Optimal Control:

- ✓ Recurrent relations for policy improvement and value iteration
- ✓ Decentralized Hilbert mapping for information fusion
- ✓ Communication protocols for efficient map-information spreading
- ✓ Information-driven roadmap results

Future Work:

- ➢ Develop adaptive DOC for effective multi-scale information gathering
- ➢ Stability analysis in the presence of delayed information propagation
- ➢ Robustness and performance analysis
- ➢ Demonstrate adaptive DOC for changing:

1. Environmental and operating conditions
2. Target conditions; 3. Platform/sensor conditions.

Acknowledgments

- **Collaborators:**

- **Pingping Zhu**, Ph.D.
  LISC, MEMS, Duke

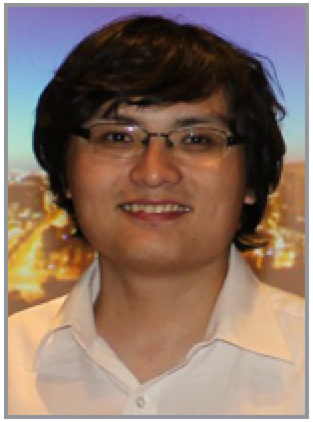

Pingping Zhu

- **Thomas A. Wettergren**, Ph.D.
  Naval Undersea Warfare Center
  Newport, RI

- **Julian Morelli**,
  LISC, MAE, Cornell

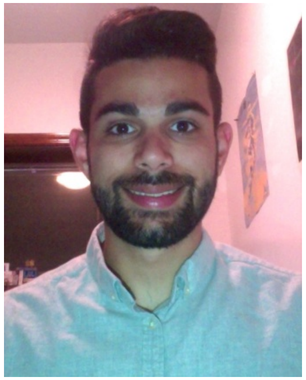

- **Collaborators:**

This work was funded by ONR Code 321

Questions?

Thank you

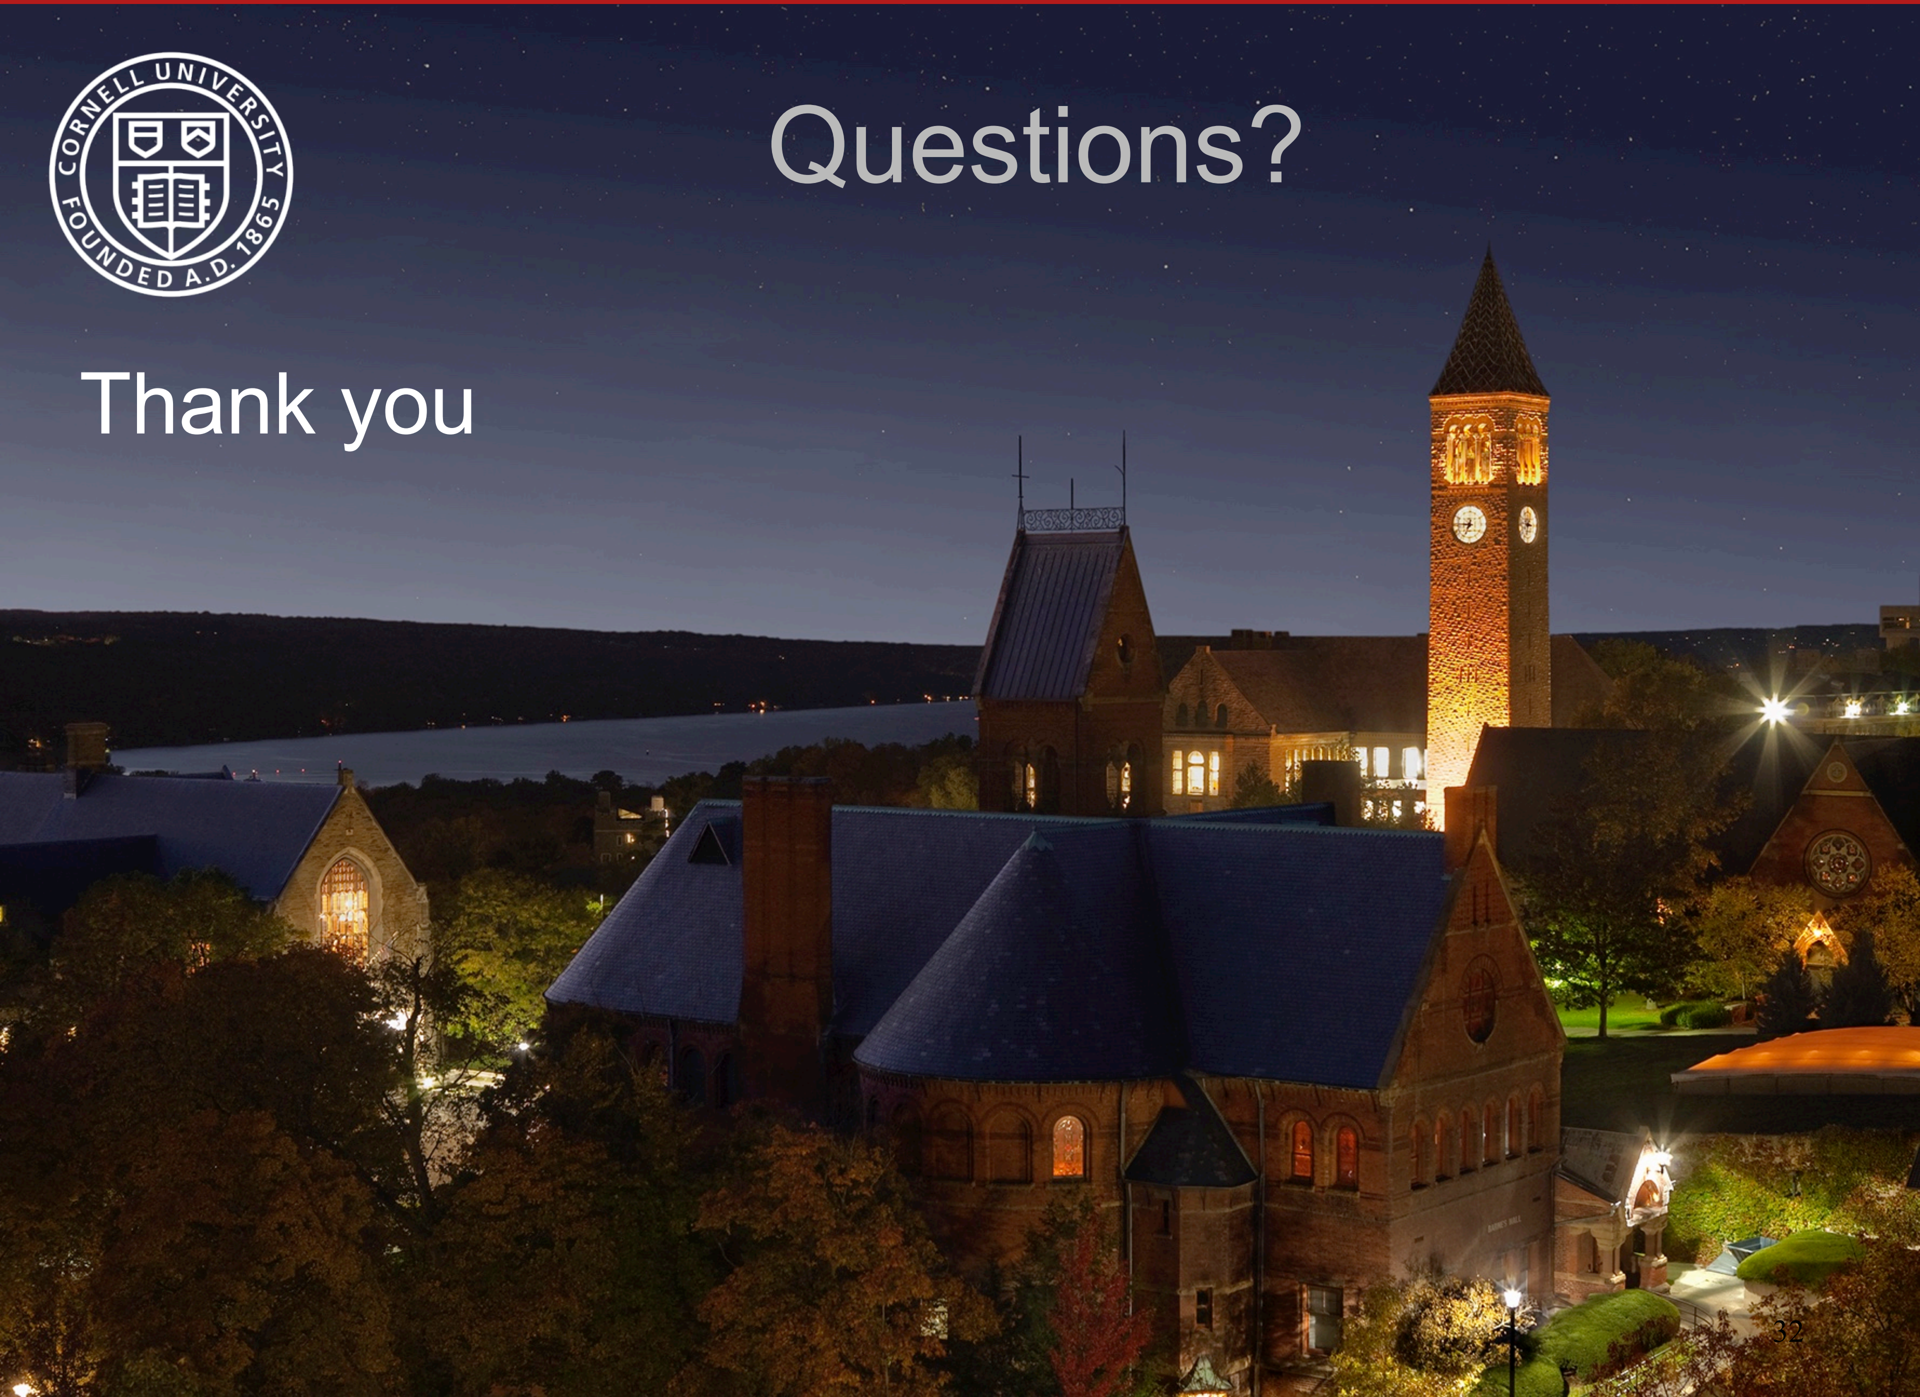

References

- K. Rudd, G. Foderaro, and S. Ferrari, “A Generalized Reduced Gradient Method for the Optimal Control of Very Large Scale Robotic (VLSR) Systems," *IEEE Transactions on Robotics*, in press.
- G. Foderaro, S. Ferrari, T. A. Wettergren, “Distributed Optimal Control of Sensor Networks for Dynamic Track Coverage," *IEEE Transactions on Control of Network Systems*, in press.
- S. Ferrari, G. Foderaro, P. Zhu, and T.Wettergren “Distributed Optimal Control of Multiscale Dynamical Systems: A Tutorial", *Control Systems Magazine*, Vol. 36, No. 2, pp. 102-116, 2016.
- G. Foderaro, S. Ferrari, T. A. Wettergren, “Distributed Optimal Control for Multi-agent Trajectory Optimization,” *Automatica*, Vol. 50, No. 1, pp. 149-154, 2014.

References

S. Ferrari and T. A. Wettergren, *Information-driven planning and Control*, CRC Press, Boca Raton, FL, scheduled to appear 2018.

S. Ferrari, G. Zhang, and C. Cai, "A Comparison of Information Functions and Search Strategies for Sensor Planning," *IEEE Transactions on Systems, Man, and Cybernetics - Part B*, in press.

N. Bezzo, R. Fierro, A. Swingler, and S. Ferrari, "Mobile Router Networks: A Disjunctive Programming Approach," *International Journal of Robotics and Automation*, in press.

S. Ferrari and G. Daugherty, "A Q-Learning Approach to Online Unmanned Air Vehicle (UAV) for Target Detection and Classification," *Journal of Defense Modeling and Simulation*, in press.

G. Foderaro, V. Raju, and S. Ferrari, "A Model-based Approximate λ-Policy Iteration Approach to Evasive Path Planning and the Video Game Ms. Pac-Man," *Journal of Control Theory and Applications*, in press.

S. Ferrari, G. Zhang, and T. A. Wettegren, "Probabilistic Track Coverage in Cooperative Sensor Networks," *IEEE Transactions on Systems, Man, and Cybernetics – Part B*, Vol. 40, No. 6, pp.1492-1504, 2011.

W. Lu, G. Zhang, S. Ferrari, R. Fierro, and I. Palunko "An Improved Particle Filter Approach for Multiple Target Detection and Tracking," *Proc. SPIE Conference*, Orlando, FL, 2011.

S. Ferrari, G. Foderaro, and A. Tremblay "A Probability Density Function Approach to Distributed Sensors Path Planning," *Proc. IEEE Conference on Robotics and Automation*, Anchorage, AK, 2010.

S. Ferrari and G. Foderaro, "A Potential Field Approach to Finding Minimum-Exposure Paths in Wireless Sensor Networks," *Proc. IEEE Conference on Robotics and Automation*, Anchorage, AK, 2010.

A. Swingler and S. Ferrari, "A Cell Decomposition Approach to Cooperative Path Planning and Collision Avoidance," *Proc. IEEE Conference on Decision and Control*, Atlanta, GA, 2010.

References

B. Bernard and S. Ferrari, “A Geometric Transversals Approach to Track Coverage of Maneuvering Targets,” *Proc. IEEE Conference on Decision and Control*, Atlanta, GA, 2010.

G. Foderaro and S. Ferrari, “Necessary Conditions for Optimality for a Distributed Optimal Control Problem,” *Proc. IEEE Conference on Decision and Control*, Atlanta, GA, 2010.

W. Lu, G. Zhang, and S. Ferrari, “A Randomized Hybrid System Approach to Coordinated Robotic Sensor Planning,” *Proc. IEEE Conference on Decision and Control*, Atlanta, GA, 2010.

S. Ferrari and G. Daugherty, "A Q-Learning Approach to Online Unmanned Air Vehicle (UAV) for Target Detection and Classification," *Proc. SPIE Conference*, Orlando, FL, April 2010.

K. C. Baumgartner, S. Ferrari, and T. A. Wettergren, “Robust Deployment of Ocean Sensor Networks,” *IEEE Sensors Journal*, Vol. 9, No. 9, pp. 1029-1048, 2009.

K. C. Baumgartner, S. Ferrari, and A. Rao, “Optimal Control of a Mobile Sensor Network for Cooperative Target Detection,” *IEEE Journal of Oceanic Engineering*, Vol. 34, No. 4, pp. 678-697, 2009.

G. Zhang, S. Ferrari, and M. Qian, “Information Roadmap Method for Robotic Sensor Path Planning,” *Journal of Intelligent and Robotic Systems*, Vol. 56, pp. 69-98, 2009.

S. Ferrari, R. Fierro, B. Perteet, C. Cai, and K. C. Baumgartner, “A Multi-Objective Optimization Approach to Detecting and Intercepting Dynamic Targets using Mobile Sensors, ” *SIAM Journal on Control and Optimization*, Vol. 48, No. 1, pp. 292-320, 2009.

S. Ferrari and C. Cai, “Information-Driven Search Strategies in the Board Game of CLUE®,” *IEEE Transactions on Systems, Man, and Cybernetics - Part B*, Vol. 39, No. 3, pp. 607-625, June 2009.

C. Cai and S. Ferrari, “Information-Driven Sensor Path Planning by Approximate Cell Decomposition,” *IEEE Transactions on Systems, Man, and Cybernetics - Part B*, Vol. 39, No. 3, pp. 672-689, June 2009.

References

S. Ferrari, “Multi-Objective Algebraic Synthesis of Neural Control Systems by Implicit Model Following,” *IEEE Transactions on Neural Networks*, Vol. 20, No. 3, pp. 406-419, March 2009.

S. Ferrari, R. Fierro, and D. Tolic, "A Geometric Optimization Approach to Tracking Maneuvering Tracking Using a Heterogeneous Mobile Sensor Network," *Proc. IEEE Conference on Decision and Control*, Shanghai, China, December 2009.

G. Zhang and S. Ferrari, "An Adaptive Artificial Potential Function Approach for Geometric Sensing," *Proc. IEEE Conference on Decision and Control*, Shanghai, China, December 2009.

G. Di Muro and S. Ferrari, "Penalty Function Method for Exploratory Adaptive-Critic Neural Network Control," *Proc. Mediterranean Conference on Control and Automation (MED'09)*, Thessaloniki, Greece, January 2009, pp. 1410-1414.

D. Tolic, R. Fierro, and S. Ferrari, "Cooperative multi-target tracking via hybrid modeling and geo- metric optimization," *Proc. Mediterranean Conference on Control and Automation (MED'09)*, Thessaloniki, Greece, January 2009, pp. 440-445.

Ferrari S., Steck J. E, and Chandramohan R., “Adaptive Feedback Control by Constrained Approximate Dynamic Programming,” *IEEE Transactions on Systems, Man, and Cybernetics - Part B: Cybernetics*, Vol. 38, No. 4, pp. 982-987, August 2008.

Baumgartner, K. C., and Ferrari, S., “A Geometric Approach to Analyzing Track Coverage in Sensor Networks”, *IEEE Transactions on Computer*, Vol. 57, No. 8, pp. 1113-1128, August 2008.

S. Ferrari and A. Vaghi, “Demining Sensor Modeling and Feature-level Fusion by Bayesian Networks,” *IEEE Sensors Journal*, Vol. 6, No. 2, pp. 471-483, April 2006.

PDFs AVAILABLE UPON REQUEST: ferrari@cornell.edu