
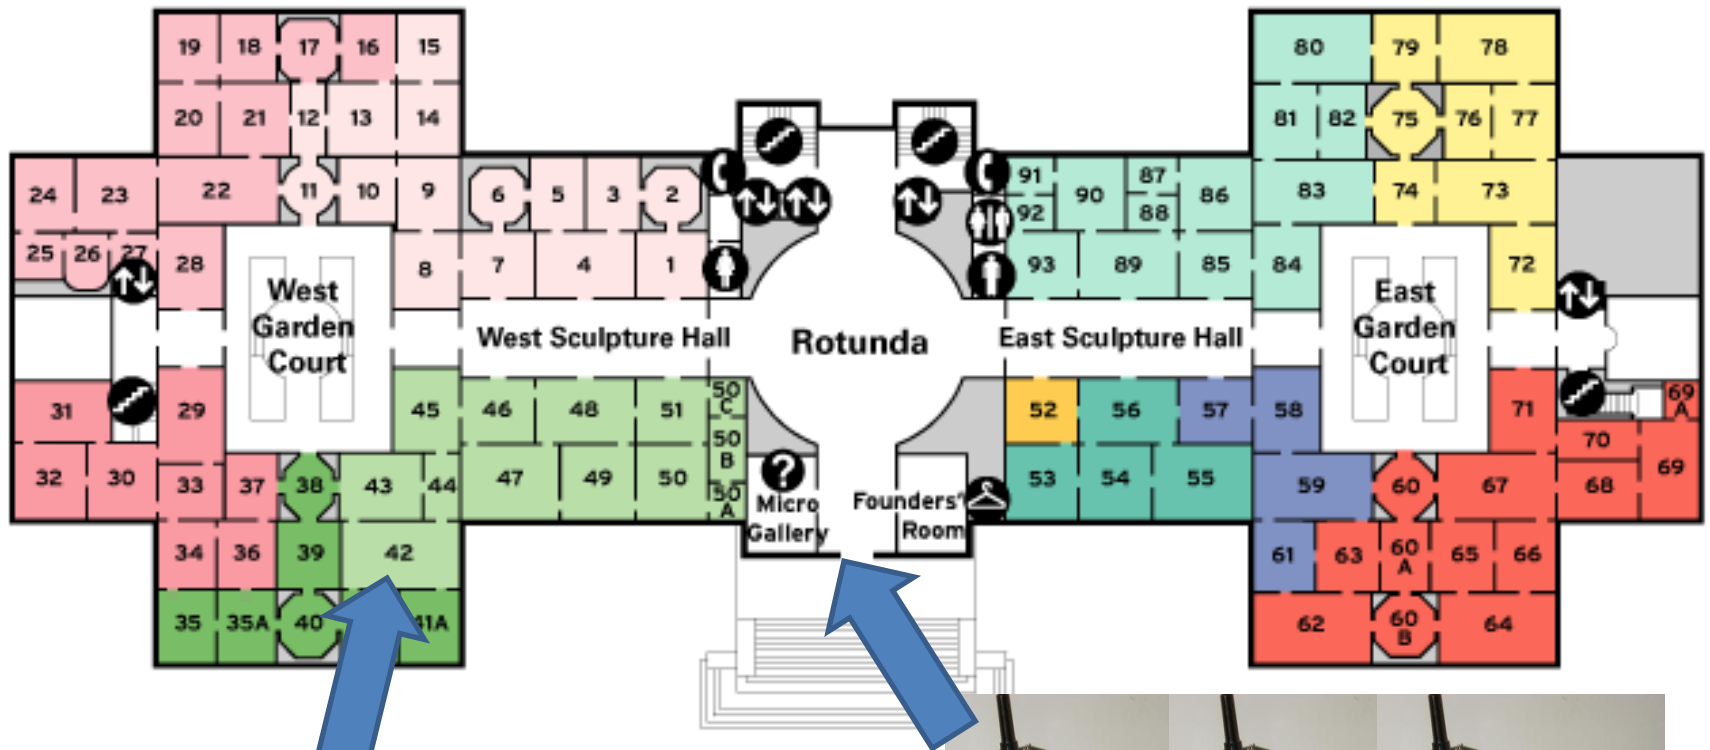


Efficient Multi-robot Search for a Moving Target

Geoffrey Hollinger, Sanjiv Singh, Joseph Djugash, Athanasios Kehagias


Christian Pothast

The problem (in pictures)



Problem: efficient search

- Quickly locate a non-adversarial target using:
 - Prior knowledge
 - A rough model of how the target moves
- Maximize **average-case** performance




(Sarmiento et al. '04, Lau et al. '06, Kurniawati et al. '09)

(Hollinger et al. IJRR '09)

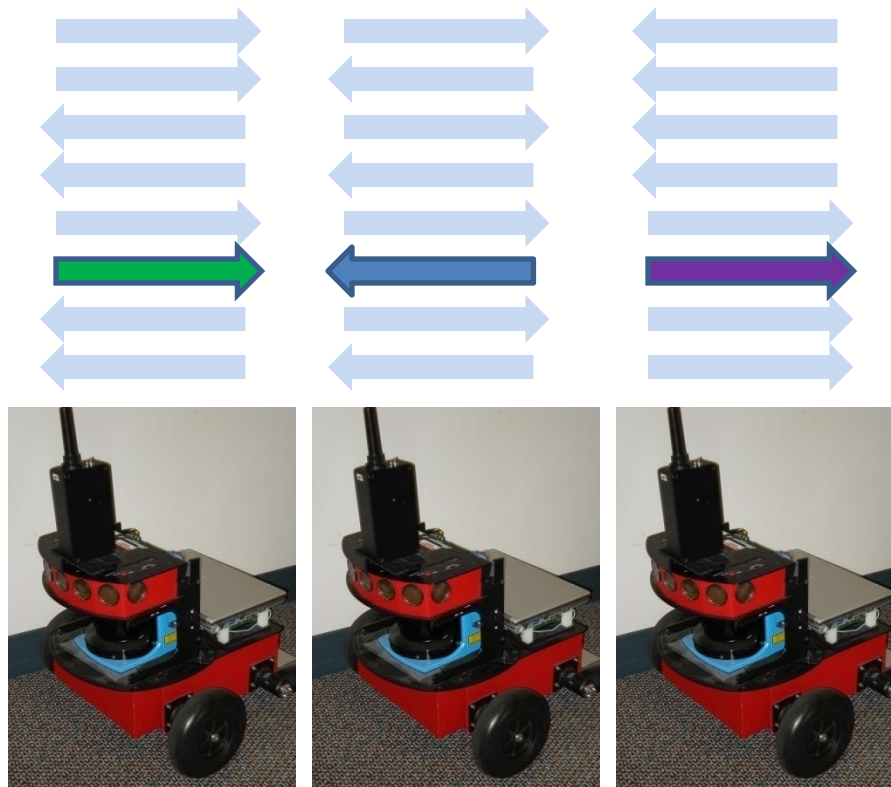
Why are multi-robot search problems hard?

- Computation scales exponentially with increasing searchers
 - Requires coordination between searchers
 - Joint action-space is the cross product of searcher actions




Why are multi-robot search problems hard?

- Computation scales exponentially with increasing searchers
 - Requires coordination between searchers
 - Joint action-space is the cross product of searcher actions



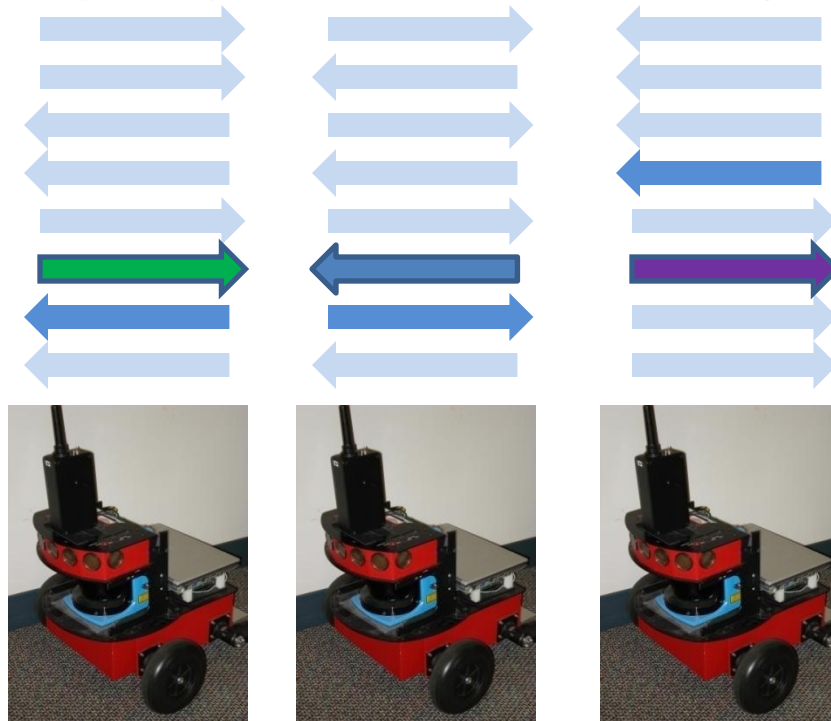
Why are multi-robot search problems hard?

- Computation scales exponentially with increasing searchers
 - Requires coordination between searchers
 - Joint action-space is the cross product of searcher actions
- Dealing with uncertainty
 - Target estimation
 - Noisy sensors



How can we tractably solve the search problem?

- Implicit coordination
 - Each robot plans *only* its own actions
 - Robots communicate information to improve their actions
 - Examples: plans, measurements, target estimation




Some quick assumptions

- Assumptions: robot capabilities
 - Robots have a coarse estimate of their global location
 - Robots can communicate locations and paths
 - Environment is known
- Assumptions: efficient search
 - Robots have line-of-sight “capture sensor”
 - May have non-line-of-sight sensors
 - Target motion is non-adversarial



Algorithm description

- The problem: maximize discounted probability of finding a target




Algorithm description

- Environment representation
 - Discretize environment into convex regions



Formulation: modeling reward

- Discounted reward metric



Formulation: modeling the target

- Discounted reward metric $F_Y(A) = \gamma^t$
- Probabilistic motion model (**average-case**)

$$\operatorname{argmax}_A \sum_{Y \in \Psi} \Pr(Y) F_Y(A)$$

- Adversarial model (**worst-case**)


$$\operatorname{argmax}_A \min_Y F_Y(A)$$

Algorithm description

- Target motion modeling and estimation
 - Define probability distribution over target's state
 - Use matrices to model capture and (Markov) diffusion


$$b(t + 1) = b(t)D \left[\prod_{k=1}^K C_{p_k(t)} \right]$$

$$D = \begin{matrix} & \begin{matrix} C & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \end{matrix} \\ \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \frac{1}{3} & 0 & 0 & \frac{1}{3} & 0 & 0 & 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{3} & 0 & 0 & 0 & 0 & 0 & \frac{1}{3} & \frac{1}{3} \\ 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 & 0 & 0 \\ 0 & \frac{1}{6} & 0 & \frac{1}{6} & \frac{1}{6} & \frac{1}{6} & \frac{1}{6} & 0 & 0 & \frac{1}{6} \\ 0 & 0 & 0 & 0 & \frac{1}{3} & \frac{1}{3} & 0 & \frac{1}{3} & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{1}{3} & 0 & \frac{1}{3} & \frac{1}{3} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 & 0 \\ 0 & \frac{1}{3} & \frac{1}{3} & 0 & 0 & 0 & 0 & 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{3} & 0 & \frac{1}{3} & 0 & 0 & 0 & 0 & \frac{1}{3} \end{pmatrix} & \begin{matrix} C \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \\ 8 \\ 9 \end{matrix} \end{matrix}$$



Algorithm description


- Maximize discounted probability of capture
- Finite-horizon path enumeration
 - Examine all paths within fixed depth
 - Replan on receding horizon
- Sequential allocation
 - For all robots
 - Robot plans using FHPE and shares path with other robots
 - Other robots assume shared paths are fixed



Pink: high probability of target
Green: medium probability of target
Blue: low probability of target

Algorithm description


- Maximize discounted probability of capture
- Finite-horizon path enumeration
 - Online
- Sequential allocation
 - Decentralized



Pink: high probability of target
Green: medium probability of target
Blue: low probability of target

Algorithm description

- Maximize discounted probability of capture
- The worst case for sequential allocation
 - Is bounded w.r.t. optimal
 - If the objective function is submodular and nondecreasing (Singh et al. '07)





Pink: high probability of target
Green: medium probability of target
Blue: low probability of target

Theoretical analysis


- Given a set function (e.g. discounted capture probability along a path)
- Nondecreasing if:
 $A \subseteq B \Rightarrow F(A) \leq F(B)$
- Submodular if:

$$A \subseteq B \Rightarrow F(A \cup C) - F(A) \geq F(B \cup C) - F(B)$$




Pink: high probability of target
Green: medium probability of target
Blue: low probability of target

Results




- Newell-Simon Hall: few cycles

- National Gallery: many cycles



Results



Simulated Office

One searcher (blue)

Randomly moving target (red)



Darker cells are more probable target locations

Results

- Our method
 - Finite-horizon path enumeration with sequential allocation
- Alternative methods
 - POMDP (optimal solution)
 - Finite-horizon explicit coordination



Results

One searcher



*Averages over 200 trials. Error bars are one SEM.


Results



*Averages over 200 trials. Error bars are one SEM.

Results

- Incorporating online sensor data



Experiments conducted with Pioneer robot in 60m x 30m office building

Results

Efficient Multi-Robot Search for a Moving Target

Geoffrey Hollinger, Sanjiv Singh,
Joseph Djugash, Athanasios Kehagias

Robotics Institute, Carnegie Mellon University
Division of Mathematics, Aristotle University of Thessaloniki

Results

Improving the Efficiency of Clearing with Multi-Agent Teams

International Journal of Robotics Research


Video Extension 2: Search with a Human-Robot Team

Authors: Geoffrey Hollinger, Athanasios Kehagias,
and Sanjiv Singh

Affiliations: Carnegie Mellon University and
Aristotle University of Thessaloniki

Results

- Incorporating online sensor data



Summary

- Algorithm for finding a non-adversarial target
 - Finite-horizon path enumeration
 - Sequential allocation
 - Online and decentralized
- Theoretical analysis
 - Linearly scalable with increasing searchers
 - Constant factor approximation guarantee
- Experimental analysis
 - Competitive with exponential algorithms
 - Incorporates data from ranging radios