

Search in the Physical World

Geoffrey A. Hollinger

June 1st, 2010

Thesis Committee:



Sanjiv Singh (chair)

Geoff Gordon

Reid Simmons

Thanasis Kehagias, Aristotle University

Search in the Physical World



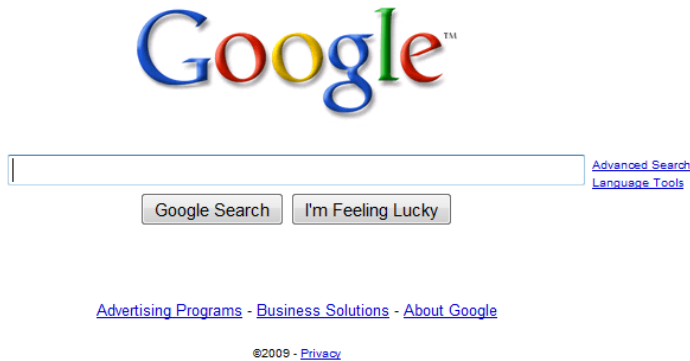
2

First Response

Reconnaissance

Wilderness Rescue

Digital world vs. Physical world



- Vast amounts of information available before search
- Much computation and memory available
- No notion of distinct cooperating agents
- **Efficient caching and recalling**
- Information gained while searching
- Limited computation and real-time operation
- Multiple searchers improve effectiveness
- **Efficient planning and coordination**


Search in the Physical world

- Design requirements
 - Multiple robots/sensors
 - Scalability
 - Computation
 - Communication
 - Decentralization
 - Robustness to failure
 - Online replanning




Why is multi-robot search hard?

- Choices increase exponentially with more searchers
 - Tasks require coordination between searchers
 - Joint action-space grows with number of searchers



Why is multi-robot search hard?

- Choices increase exponentially with more searchers
 - Tasks require coordination between searchers
 - Joint action-space grows with number of searchers

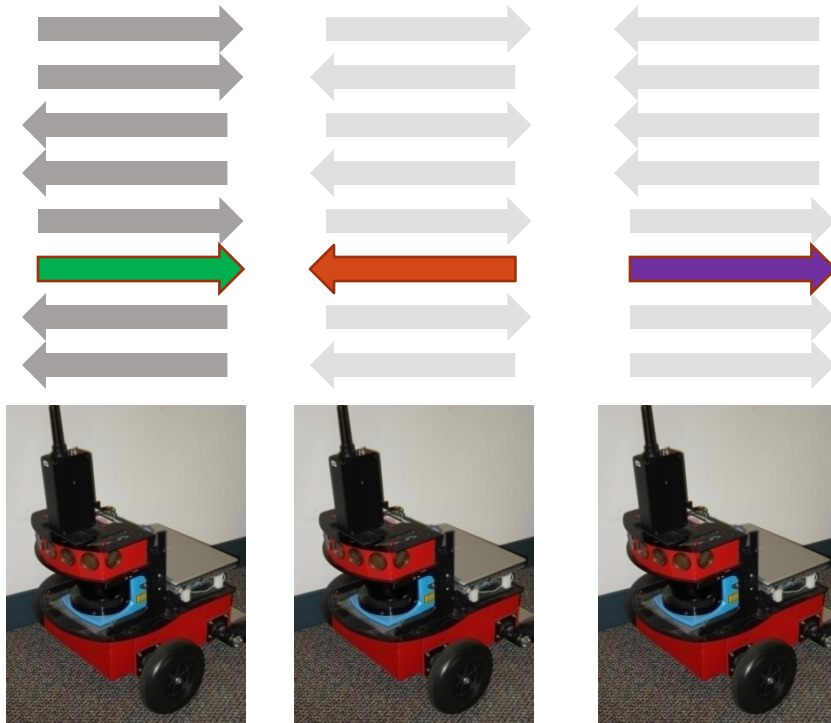


One solution

- Centralized solver (Smith et al. '05, Kurniawati et al. '08)
 - One leader plans for the entire team
 - The leader tells everyone else what to do

Advantages

- Considers all possible strategies



Disadvantages

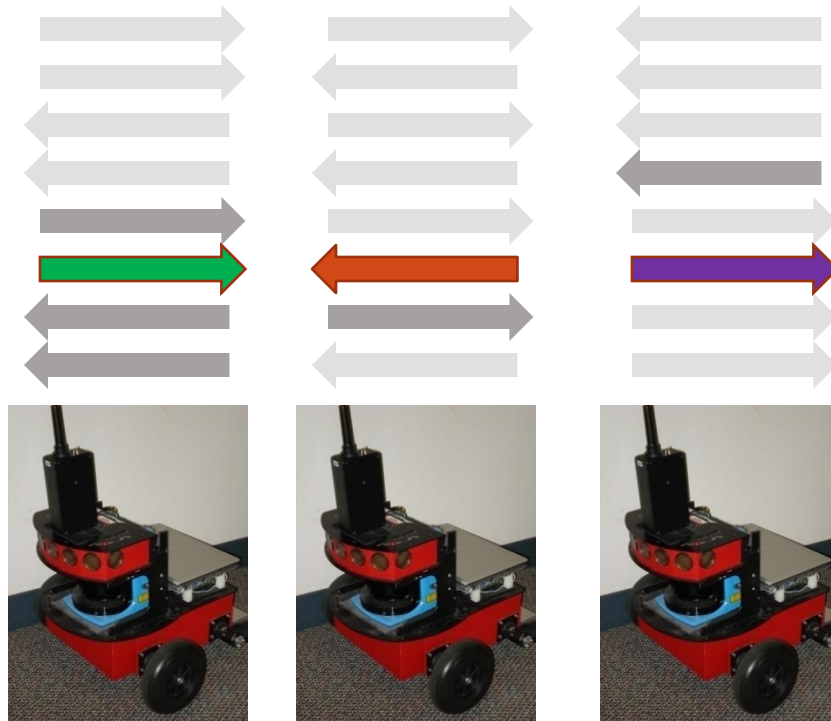
- High computation
- Single-point failure

A better solution

- Market-based coordination (Kalra et al. '06, Zlot et al. '06)
 - Each robot plans for itself and for some teammates
 - Robots auction plans to other robots

Advantages

- Relatively scalable
- Robust
- Decentralized



Disadvantages

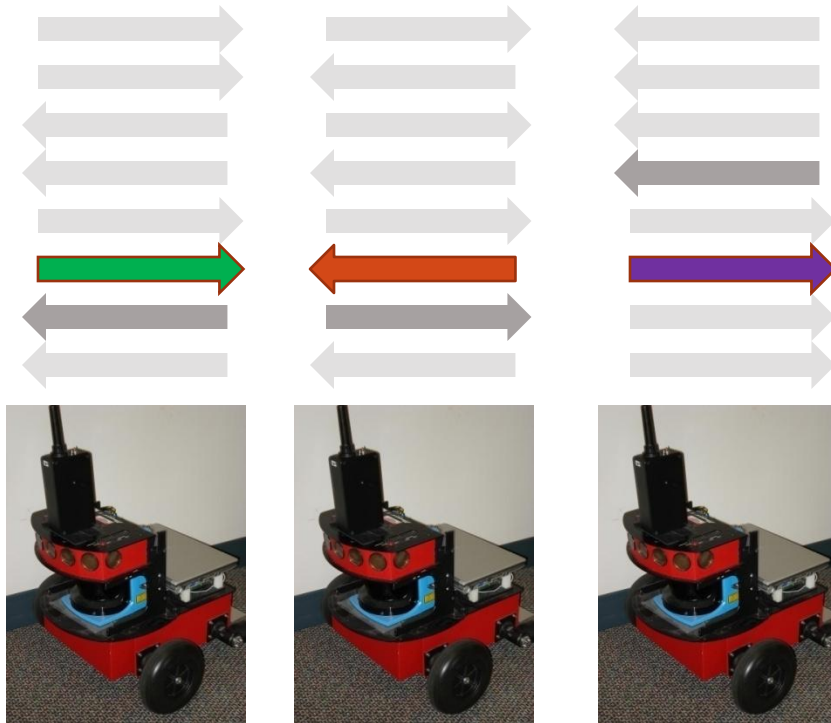
- Higher communication
- Need to determine when to auction
- Often no bounds on performance

Proposed solution

- Implicit coordination (Hollinger et al. '07-'10)
 - Each robot plans **only** its own actions
 - Robots communicate plans to improve their actions

Advantages

- Scalable
- Robust
- Decentralized
- Little communication



Disadvantages

- How well does it work?
- Can we analyze it formally?

Thesis Statement

- Algorithms using implicit coordination provide scalable, decentralized, online, and high-performing solutions to multi-robot search problems in complex physical environments.

Thesis Statement

- Algorithms using implicit coordination provide scalable, decentralized, online, and high-performing solutions to **multi-robot search** problems in complex **physical environments**.

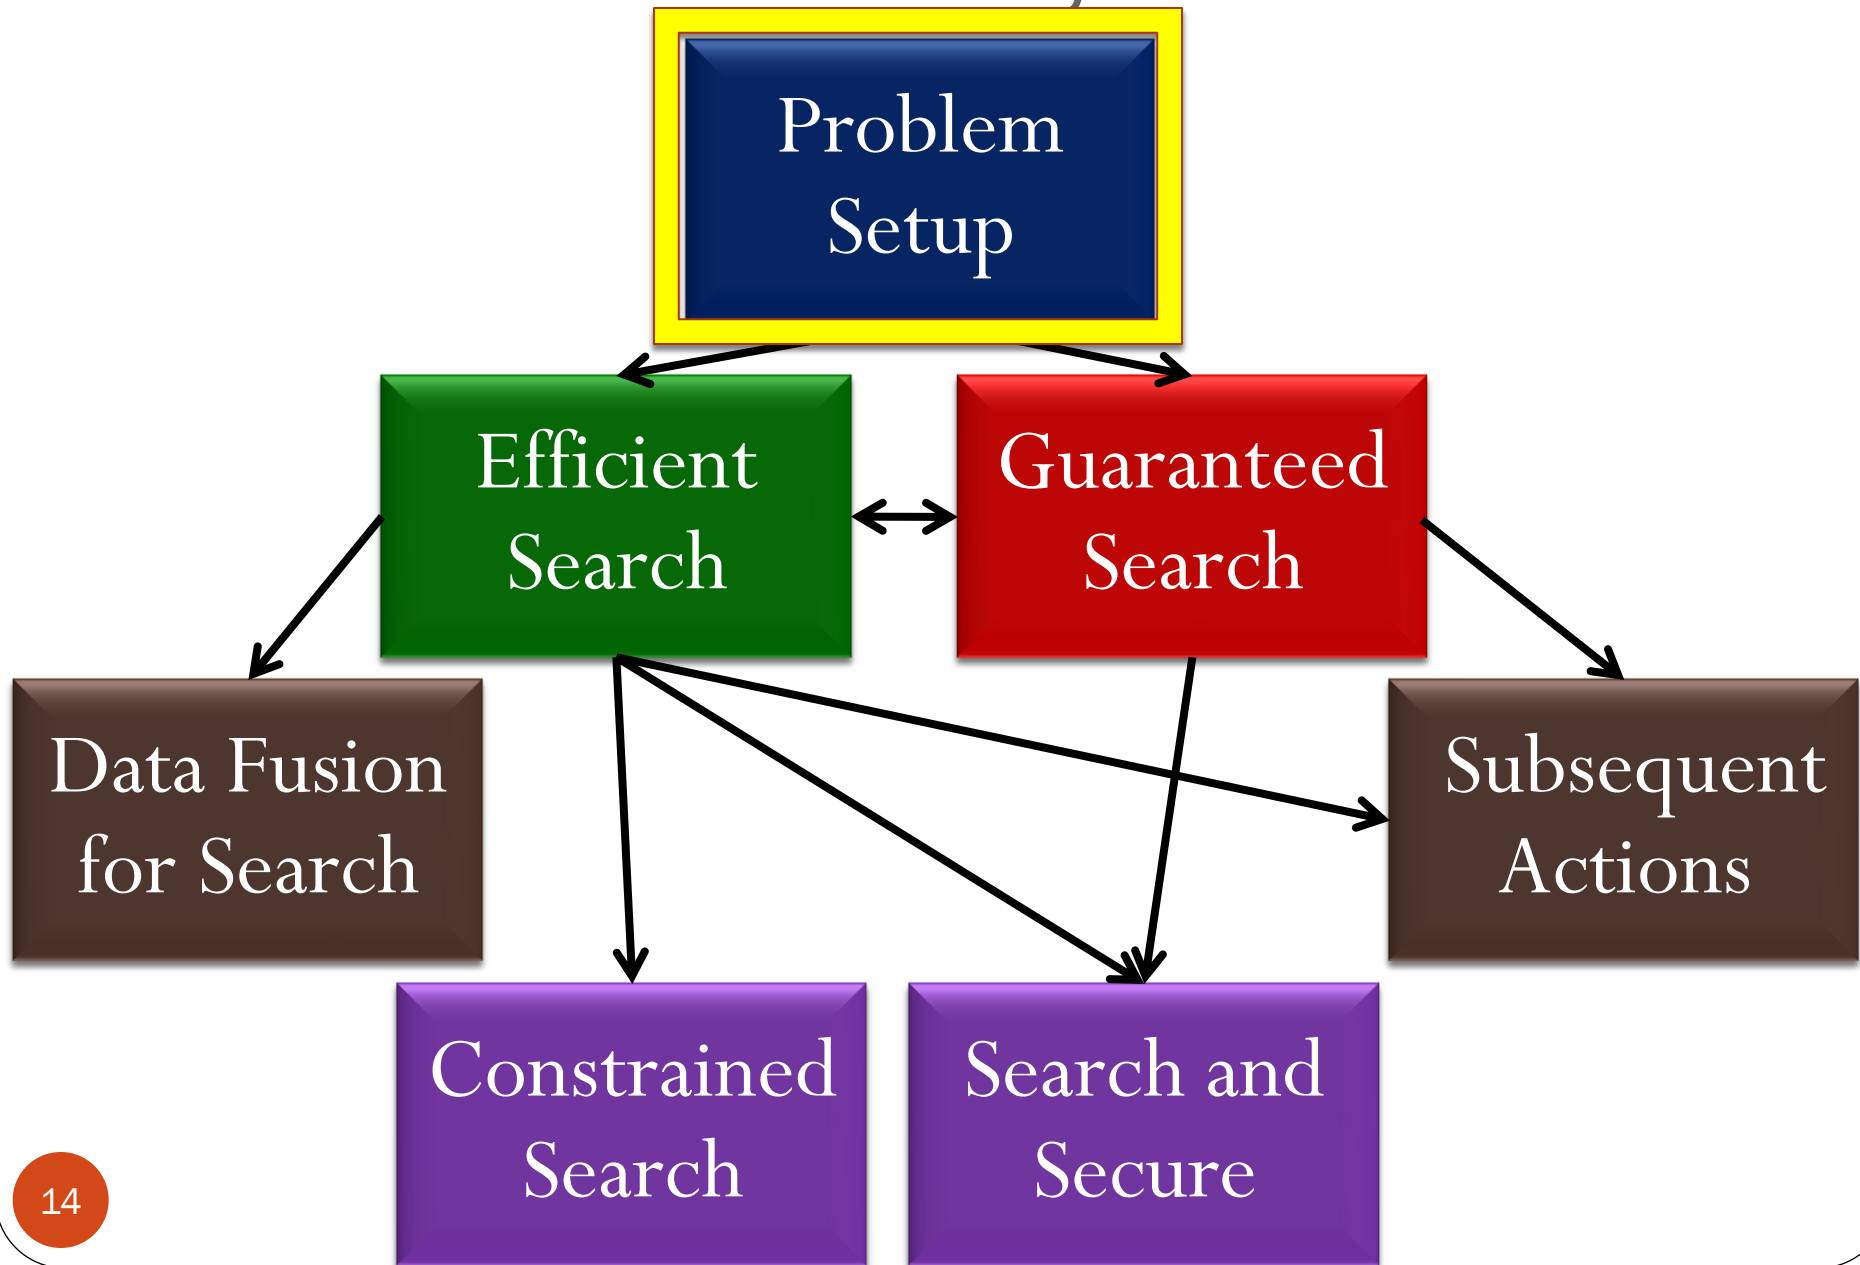
Thesis Statement

- Algorithms using implicit coordination provide **scalable, decentralized, online, and high-performing** solutions to **multi-robot search** problems in complex **physical environments**.


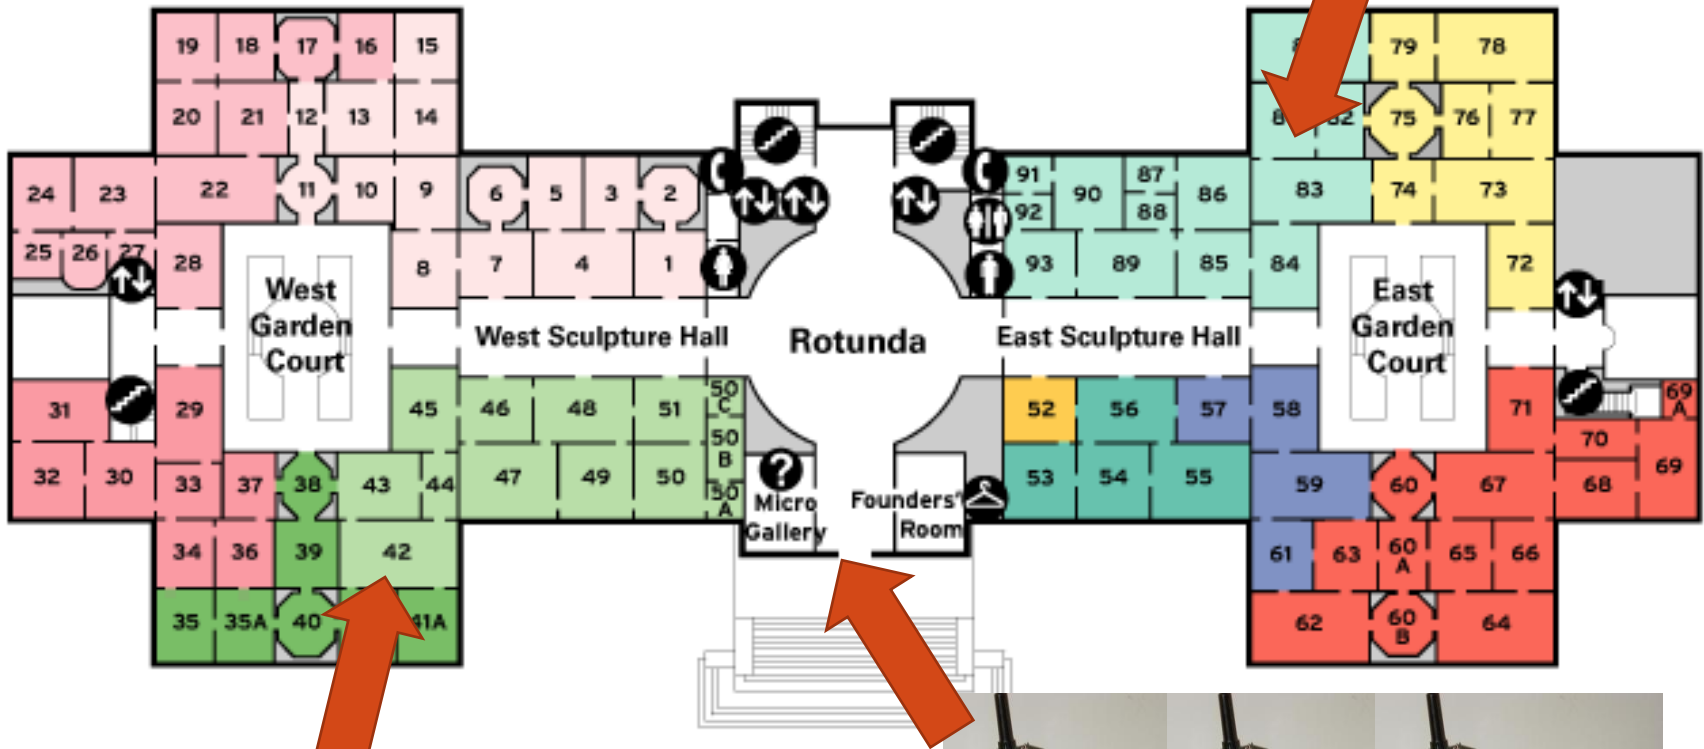
Thesis Statement

- Algorithms using **implicit coordination** provide **scalable, decentralized, online, and high-performing** solutions to **multi-robot search** problems in complex **physical environments**.

Search in the Physical World




Two search problems



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)

Problem: guaranteed search

- Guarantee clearing an environment of any potential threat
 - Adversarial target
 - Potentially infinite speed
- Optimize **worst-case** performance




(Parsons '76, LaValle et al. '97, Gerkey et al. '05)

(Hollinger et al. AURO '10)

Problem: efficient, guaranteed search


- Clear environment of any potential threat
 - Efficiency tradeoff
 - Quickly clear the environment
 - Optimize a non-adversarial model while clearing
- Combine **average-case** and **worst-case** performance



(Hollinger et al. IJRR '10)

Problem: efficient, guaranteed search


- Clear environment of any potential threat
 - Efficiency tradeoff
 - Quickly clear the environment
 - Optimize a non-adversarial model while clearing
- Combine **average-case** and **worst-case** performance



(Hollinger et al. IJRR '10)

Problem: constrained search



- Maximize **efficient** or **guaranteed** search objective
- Searcher communication is **constrained**
 - Line-of-sight, range, etc.
- Increases necessary coordination



(Hollinger et al. ICRA '10)



Formulation: representation

- Descretize environment into convex cells
 - Constrained Delaunay triangulation (Schewchuk '02)
 - Voronoi tessellations (Kolling and Karpin '08)
- Generate search graph




Formulation: representation

- Descretize environment into convex cells
 - Constrained Delaunay triangulation (Schewchuk '02)
 - Voronoi tessellations (Kolling and Karpin '08)
- Generate search graph



Formulation: modeling reward

- Discounted reward metric $F_Y(A) = \gamma^t$



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**)

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


- Combined search problem

$$\operatorname{argmax}_A \alpha \sum_{Y \in \Psi} \Pr(Y) F_Y(A) + (1 - \alpha) \min_Y F_Y(A)$$

Why is the combined problem hard?


- Both optimization problems are NP-hard
- Good, fast approximations possible for **efficient** search
(Hollinger et al. IJRR '09)
- **Guaranteed** and **constrained** search are inapproximable*
(Hollinger et al. IJRR '10)
- So what do we do?!?!

Search in the Physical World



Algorithm: optimizing over the model


- Maximize discounted probability of capture
- Finite-horizon path enumeration
 - Examine all paths within fixed depth
 - Replan on receding horizon
- Sequential allocation
 - For all searchers
 - Plan using FHPE and share path with others
 - Other searchers assume shared paths fixed
- Linear scalability in number of searchers



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

Algorithm: optimizing over the model



- 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

Theoretical analysis



- Theorem: Multi-robot efficient search path planning optimizes a nondecreasing, submodular set function.
- This means: Sequential allocation yields a constant factor approximation guarantee.



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



Efficient search: method comparison

- Proposed method
 - FHPE+SA ←
- Alternative methods
 - POMDP (Smith et al. '07)
 - Finite-horizon explicit coordination





Efficient search: results



One searcher



Efficient search: results




Efficient search: computation time




- HSVI (POMDP):
 - Unable to find a meaningful policy even with one searcher


Efficient search: movies



One searcher, expected capture: 146 steps



Two searchers, expected capture: 68 steps




Three searchers, expected capture: 45 steps

Sensed area 

Darker areas are more probable target locations




Search in the Physical World




Implicit coordination for clearing?

- Clearing with limited robots requires coordination
- In this example
 - Greedy clearing requires $N+1$ robots
 - Clearing possible with three robots





Algorithm: finding a clearing schedule

- Clearing is NP-hard on graphs
 - Linear-time on trees (Barriere et al. '02)
- Generate a spanning tree
 - Calculate labeling (optimal for tree)
 - Generate clearing schedule from labeling with implicit coordination
 - May need more searchers than the tree
- Linear scalability
- Anytime capabilities
- Probabilistically complete



Minimizing clearing time


- PARISH⁺⁺: parallel stochastic hill-climbing (Gerkey et al. '05)
- FHGC: extension of sequential allocation
- G-GSST: spanning tree guided ←





Minimizing clearing time

- PARISH⁺⁺: parallel stochastic hill-climbing (Gerkey et al. '05)
- FHGC: extension of sequential allocation
- G-GSST: spanning tree guided ←


Lower
better



40



Search in the Physical World



Combined algorithm: put them together

- Minimize the average-case capture time
 - Efficient search
 - Receding horizon sequential allocation




- Minimize the worst-case capture time
 - Guaranteed search
 - Spanning tree guided clearing




Anytime combined algorithm

- Given a search graph and K searchers
- While time left (pre-planning)
 - Generate spanning tree
 - For decreasing number of worst-case searchers
 - Generate search schedule using:
 - Some worst-case searchers (clearing)
 - The rest average-case searchers (optimizing)
 - If clearing strategy infeasible: break
- Run best strategy (execution)



Combined algorithm movie




Simulated search in an office building.
Four guaranteed searchers (magenta) clear
the environment of a worst-case target.


Combined algorithm results

- 10,000 spanning trees (~2 minutes)

Office: N = 5 Total Searchers




Museum: N = 7 Total Searchers




Combined algorithm results

- 100 spanning trees (~1 second)

Office: N = 5 Total Searchers




Museum: N = 7 Total Searchers




Combined algorithm results

- 500 spanning trees (~5 seconds)

Office: N = 5 Total Searchers



Museum: N = 7 Total Searchers




Office schedules stop improving
(~10% of schedules considered)


Combined algorithm results

- 1000 spanning trees (~10 seconds)

Office: N = 5 Total Searchers




Museum: N = 7 Total Searchers




Combined algorithm results

- 5000 spanning trees (~1 minute)

Office: N = 5 Total Searchers



Museum: N = 7 Total Searchers




Museum schedules stop improving
(«1% of schedules considered)


Combined algorithm results

- 10,000 spanning trees (~2 minutes)

Office: N = 5 Total Searchers




Museum: N = 7 Total Searchers




Results: human-robot search team

- Shared pre-processing
 - Searchers decide on underlying spanning tree
 - Efficient and guaranteed searchers assigned
- Plan execution
 - Distributed: each agent plans its own path
 - Online: agents refine plans as they go




Results: human-robot search team



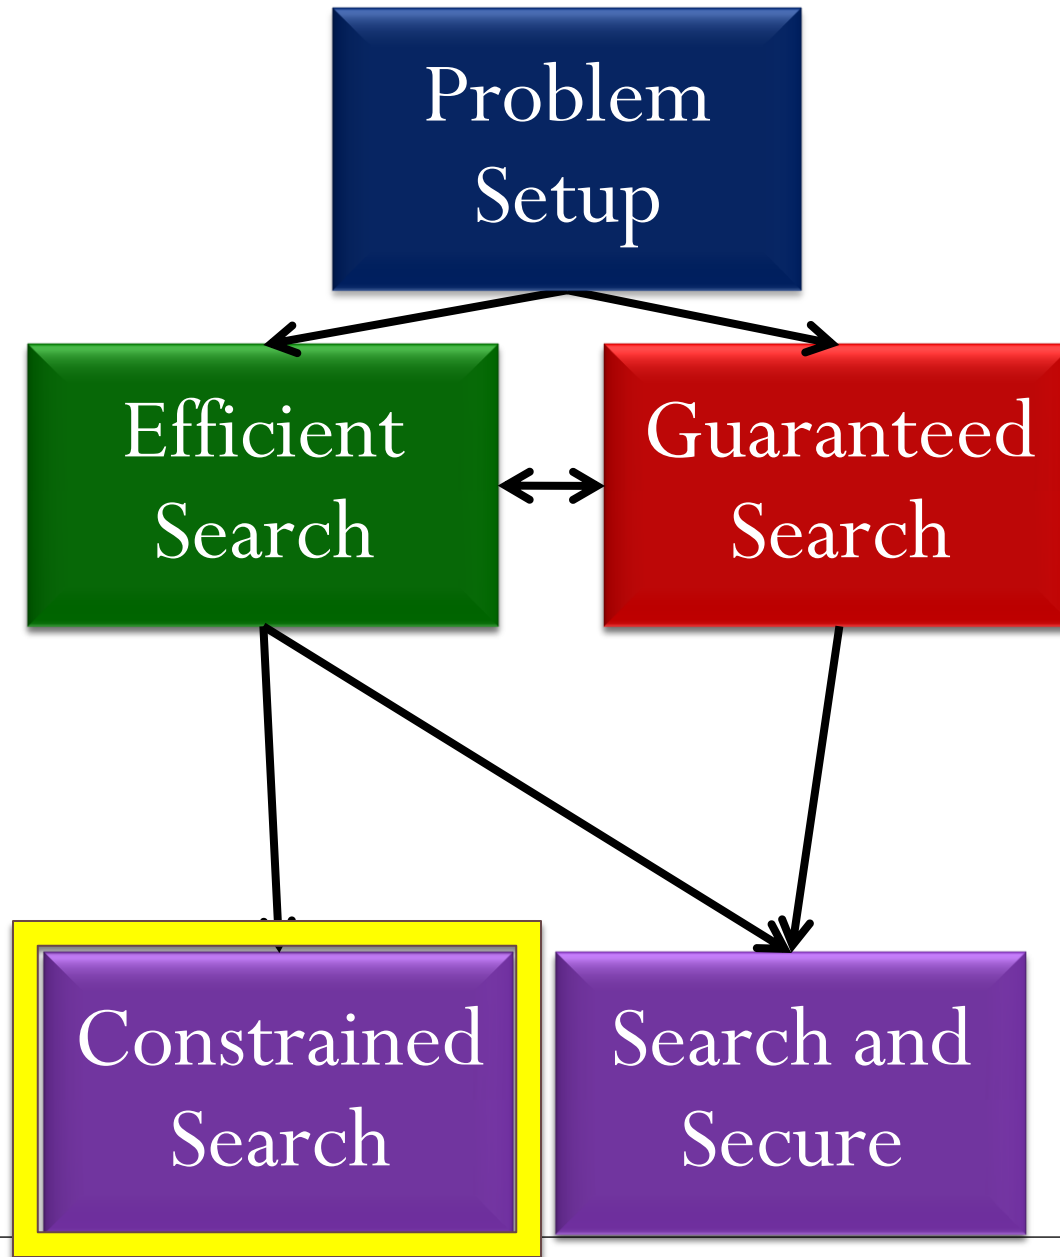
Human-robot team experiments conducted in 60 m x 30 m office building.

Results: human-robot search team







Instead, Pioneer optimizes over a target motion model. Location still shared with humans.

Search in the Physical World




Constrained search

We need to search this environment, but our communication is restricted by range and line-of-sight.




A common approach

If we take these paths we can stay in communication.



Proposed approach

If we take these shorter paths we can regain communication on the other side.



A spectrum of connectivity



Continual connectivity

(Zavlanos et al. '08, Hsieh et al. '08)

- Advantages
 - More robust
 - Online capabilities
 - Easier data fusion
- Disadvantages
 - Limited path choices
 - No performance guarantees

No connectivity

(Singh et al. '09)

- Advantages
 - Unlimited path choices
 - Known performance guarantees
- Disadvantages
 - Less robust
 - Offline operation
 - Harder data fusion

Connectivity constrained planning

- Given a discrete representation of the environment
- Need to optimize some function of K robots' paths


$$\operatorname{argmax}_{A_1, \dots, A_K} F(A_1, \dots, A_K)$$

- And regain connectivity every T_i steps

Why is this problem hard?

- Path optimization with connectivity is inapproximable*
 - Even when the underlying problem is not
- Why?
 - Finding a connected path from A to B is hard (NP-complete)


Can we get from start to goal and maintain range constraints?



Why is this problem hard?

- Path optimization with connectivity is inapproximable*
 - Even when the underlying problem is not
- Why?
 - Finding a connected path from A to B is hard (NP-complete)
 - Optimizing over connected paths is really hard (inapproximable)


What's the best path that maintains constraints *and* maximizes the probability of finding someone?



Implicit coordination algorithm


- What if we allow for periodic connectivity and apply implicit coordination?

Implicit coordination algorithm




- While not converged
 - For all robots
 - Enumerate some paths up to time T_i
 - Discard paths disconnected from team at time T_i
 - Choose path that best optimizes objective
 - Share path
 - Execute path and regain connectivity at time T_i
 - Repeat

Implicit coordination algorithm



- Scalable
- Decentralized
- Online replanning
- Converges
(if properly synchronized)
- Performance guarantees
(as periodic interval increases)

Constrained search: movie



Four searchers, simulated outdoor environment



Randomly moving target

Darker cells are more probable target locations

Objectives:

- 1) Minimize expected capture time
- 2) Maintain periodic line-of-sight connectivity

Constrained search: movie




Coupled gradient
Continual connectivity
Expected capture time:
66 steps

Implicit coordination
Periodic connectivity
Expected capture time:
45 steps

Capture time reduced by 32%

Constrained search: movie



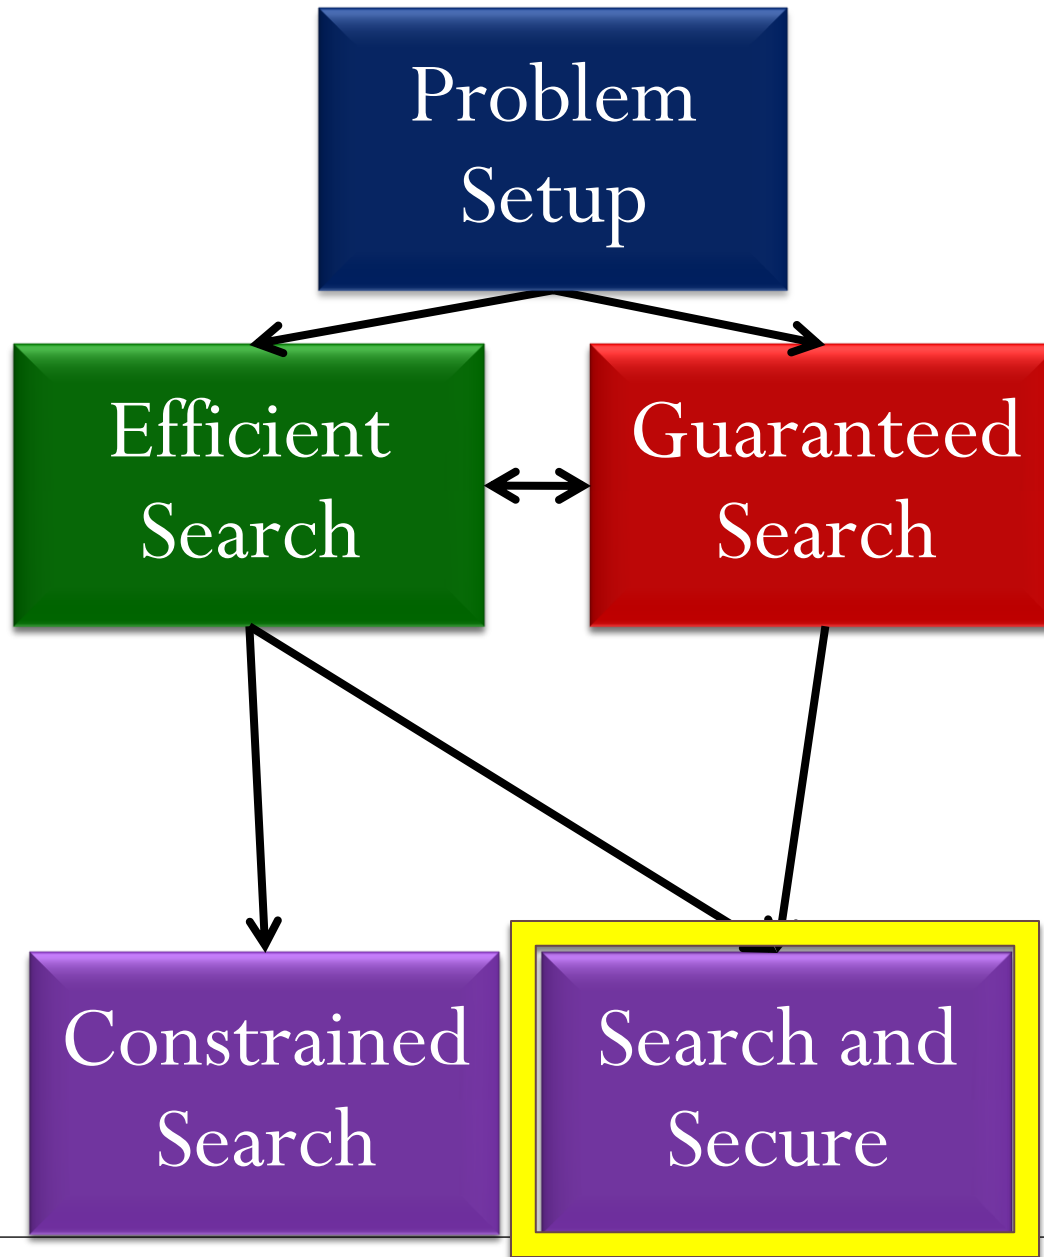
Sideswipe and Serf: differential drive robots

Experiments performed in an outdoor patio

Sensors: laser scanner (coverage),
ranging radios (localization)

Onboard computation (localization and path planning)


Search in the Physical World



MAGIC 2010

- International competition
- Search and secure
 - Find and secure all targets

Moving Target



Disruptor


Searcher

MAGIC 2010

- International competition
- Search and secure
 - Find and secure all targets



Moving Target




Disruptor

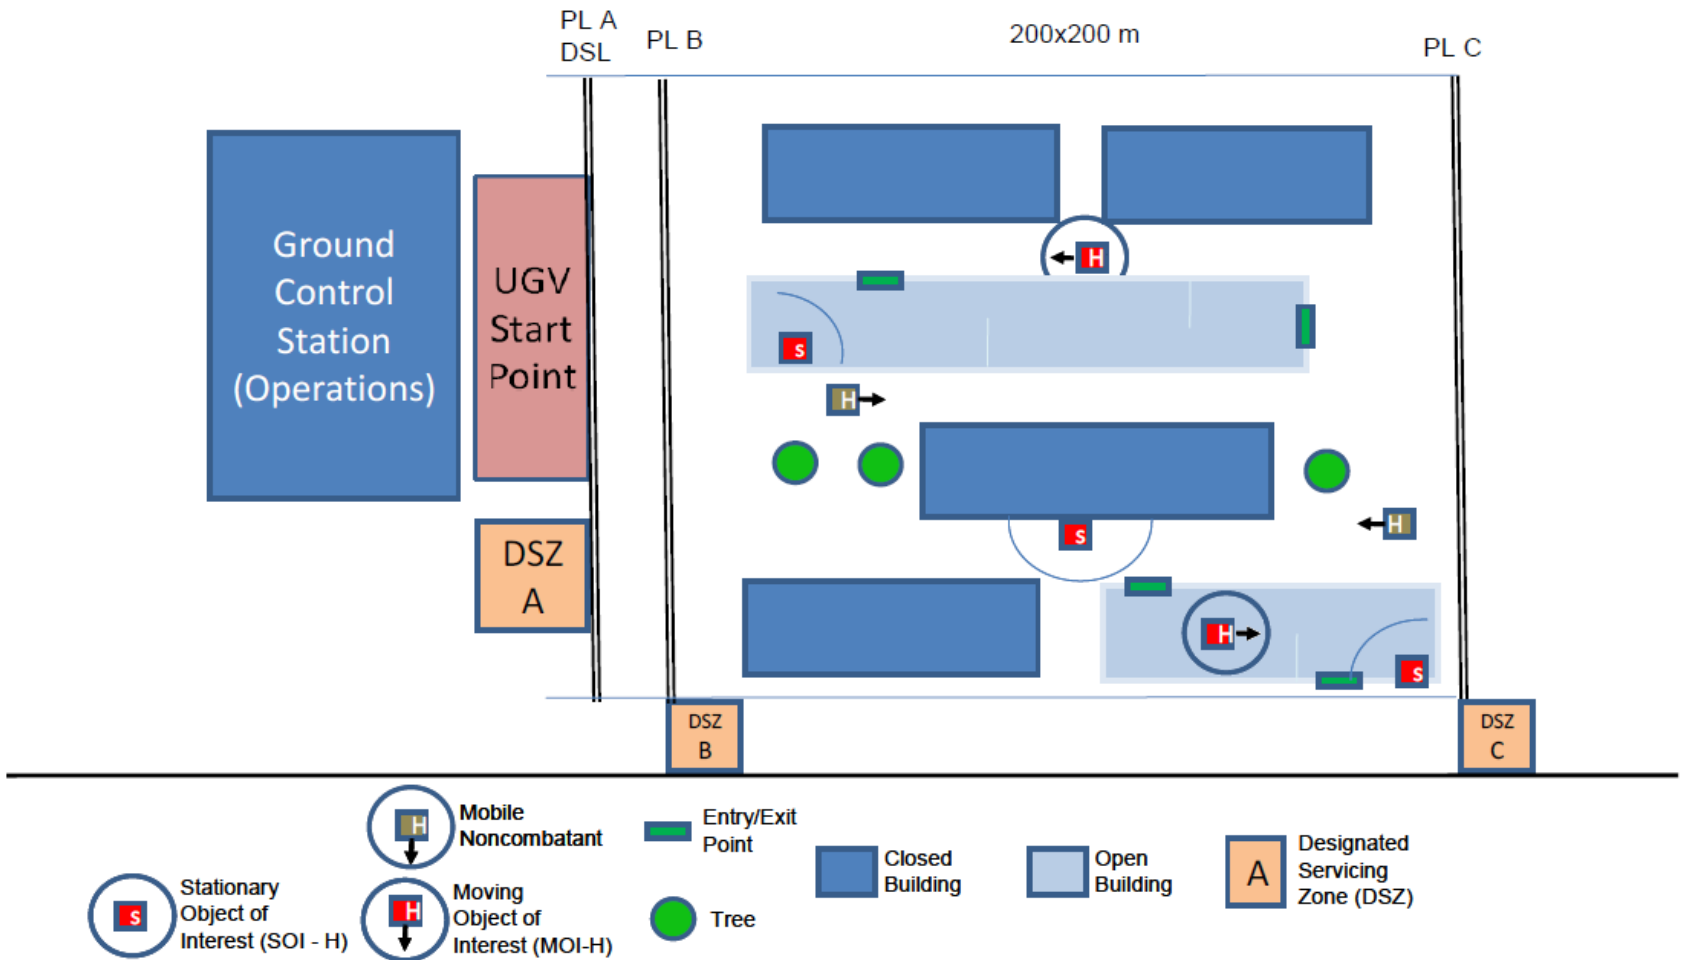
Searcher

MAGIC 2010

- International competition
- Search and secure
 - Find and secure all targets





MAGIC 2010 scenario




MAGIC 2010 solution

- Discretize using a region growing technique



- Find five searcher clearing schedule with G-GSST
- Coordinate disruptors using sequential allocation
- Replanning and schedule repair





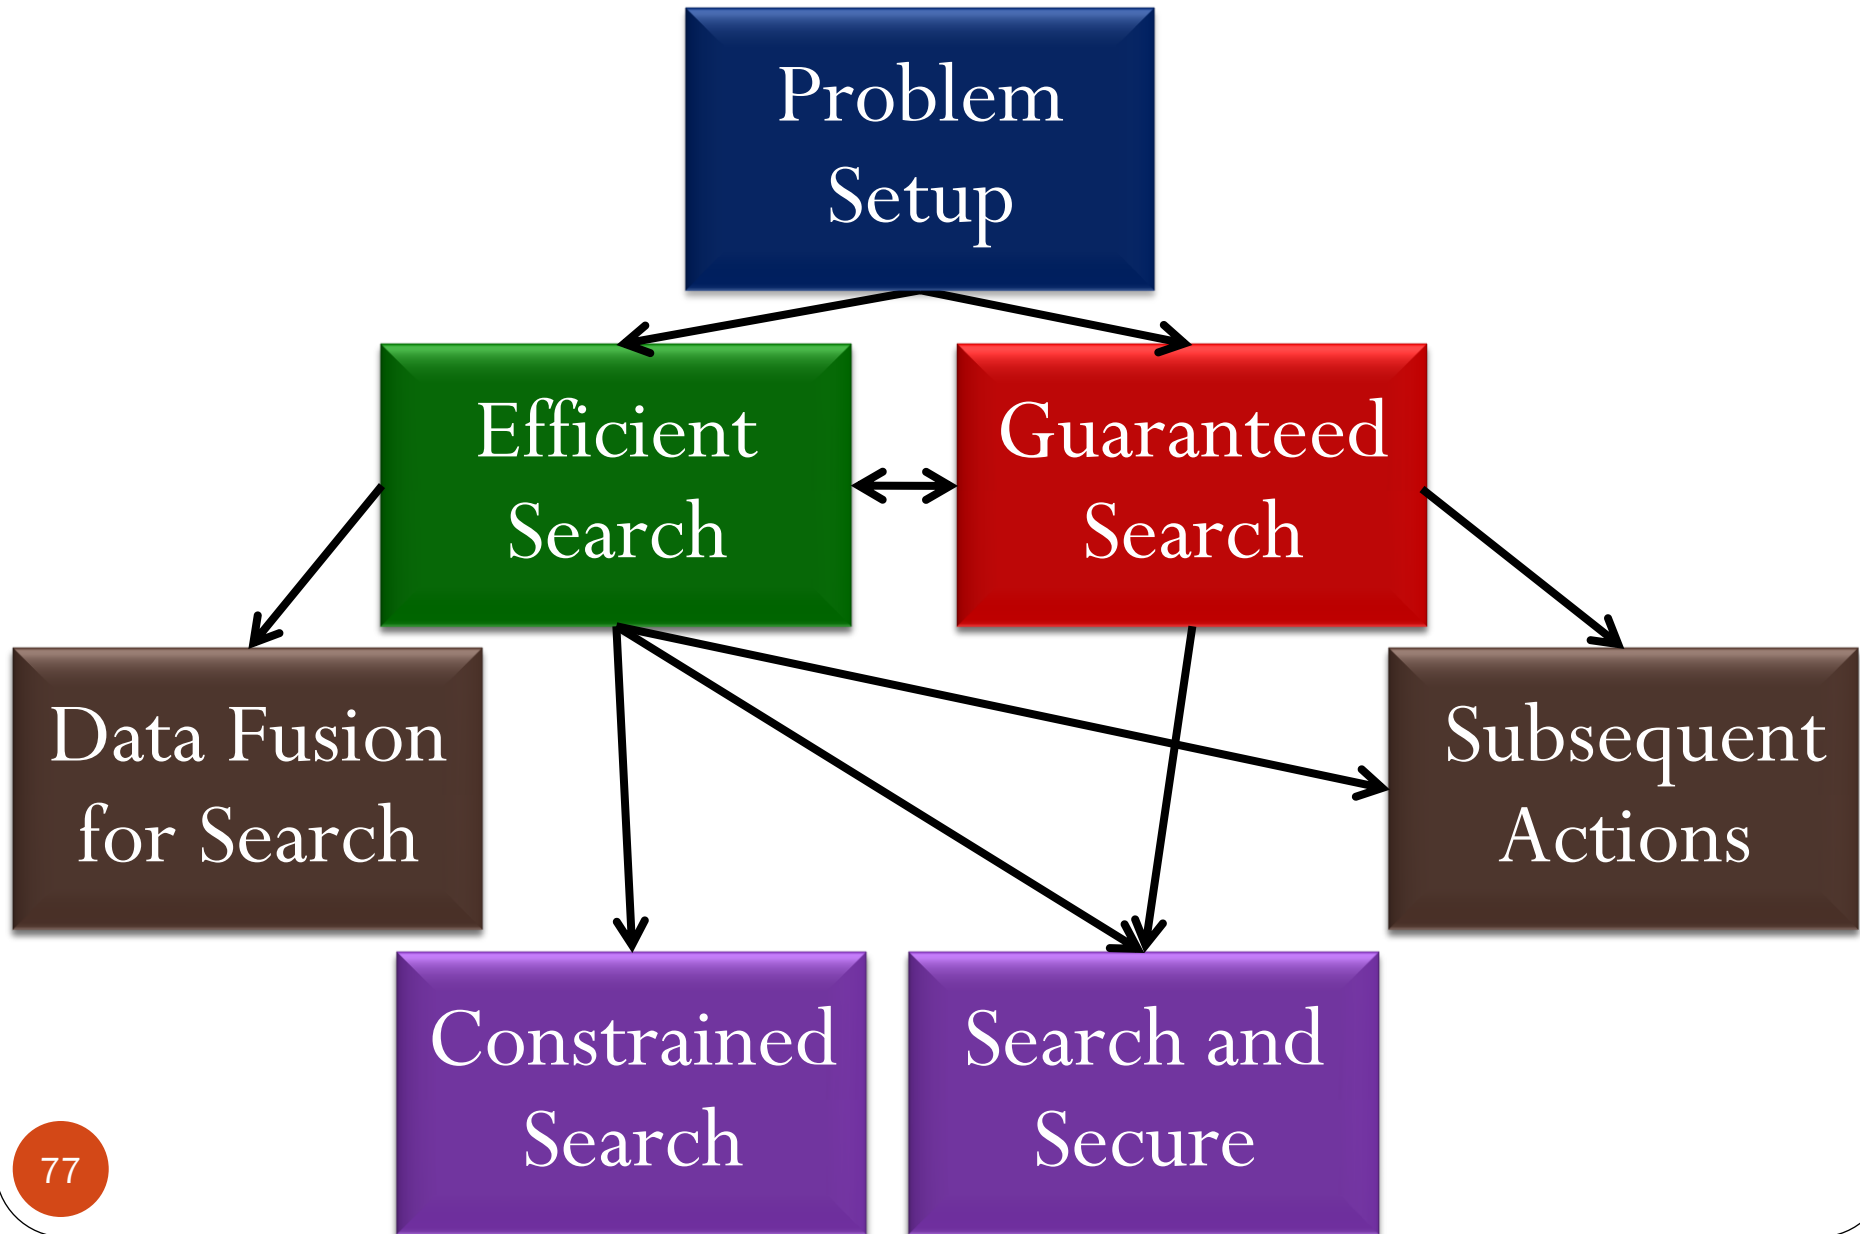
Simulated MAGIC results

Averages over 200 trials

5 Searchers, 2 Disruptors	Percentage	Avg. Completion Time
Completion (no robots disabled)	72.5%	687.41 ± 39.48 sec
Completion (robot disabled, clearing repaired)	8.0%	755.69 ± 57.75 sec
Completion (robot disabled, resort to efficient)	9.5%	854.42 ± 313.84 sec
----- Unable to complete (all disruptors disabled)	10.0%	-

- High success rates (90%)
- Failure recovery often possible
- Efficient search recoveries are “risky”


Search in the Physical World



Contributions

- Formalization of multi-robot search
- First integrated framework for efficient, guaranteed, and constrained search that is:
 - Scalable
 - Decentralized
 - Online (replanning)
 - High-performing (outperforms prior techniques)
- Validation in simulation and on human/robot teams
- Application to MAGIC 2010 search and secure problem

One last take home “algorithm”

- Given a problem that requires multiple robots
 - Does the objective have diminishing returns?
 - Can you find an informed representation?
 - Can you relax some inter-robot constraints?
- If yes to one or more  apply implicit coordination



Thanks!


- My advisor: Sanjiv Singh
- Thesis committee: Thanasis Kehagias, Geoff Gordon, Reid Simmons
- Colleagues: Joe Djugash, Ben Grocholsky, Brad Hamner, Pragun Goyle, Sidd Srinivasa, Dave Ferguson, Chris Geyer, Debadepta Dey, Matt Aasted
- Suzanne Lyons Muth and Sanae Minick
- Friends and family


- NSF Grant No. IIS-0426945, The Boeing Company, Intel Research Pittsburgh, The Robotics Institute

Thanks!

Questions?

gholling@ri.cmu.edu





- What do you think, four spiders?
- Let's do eight, I gotta get home to eat.