Planning and Inference for Micro-air Vehicle Flight in GPS-Denied Environments

Nicholas Roy

Tuesday, June 1st, 2010

AERO ASTRO

MAV 08

1ST US-ASIAN DEMONSTRATION & ASSESSMENT OF MICRO-AERIAL & UNMANNED GROUND VEHICLE TECHNOLOGY

10-15 MARCH, 2008 | AGRA, INDIA

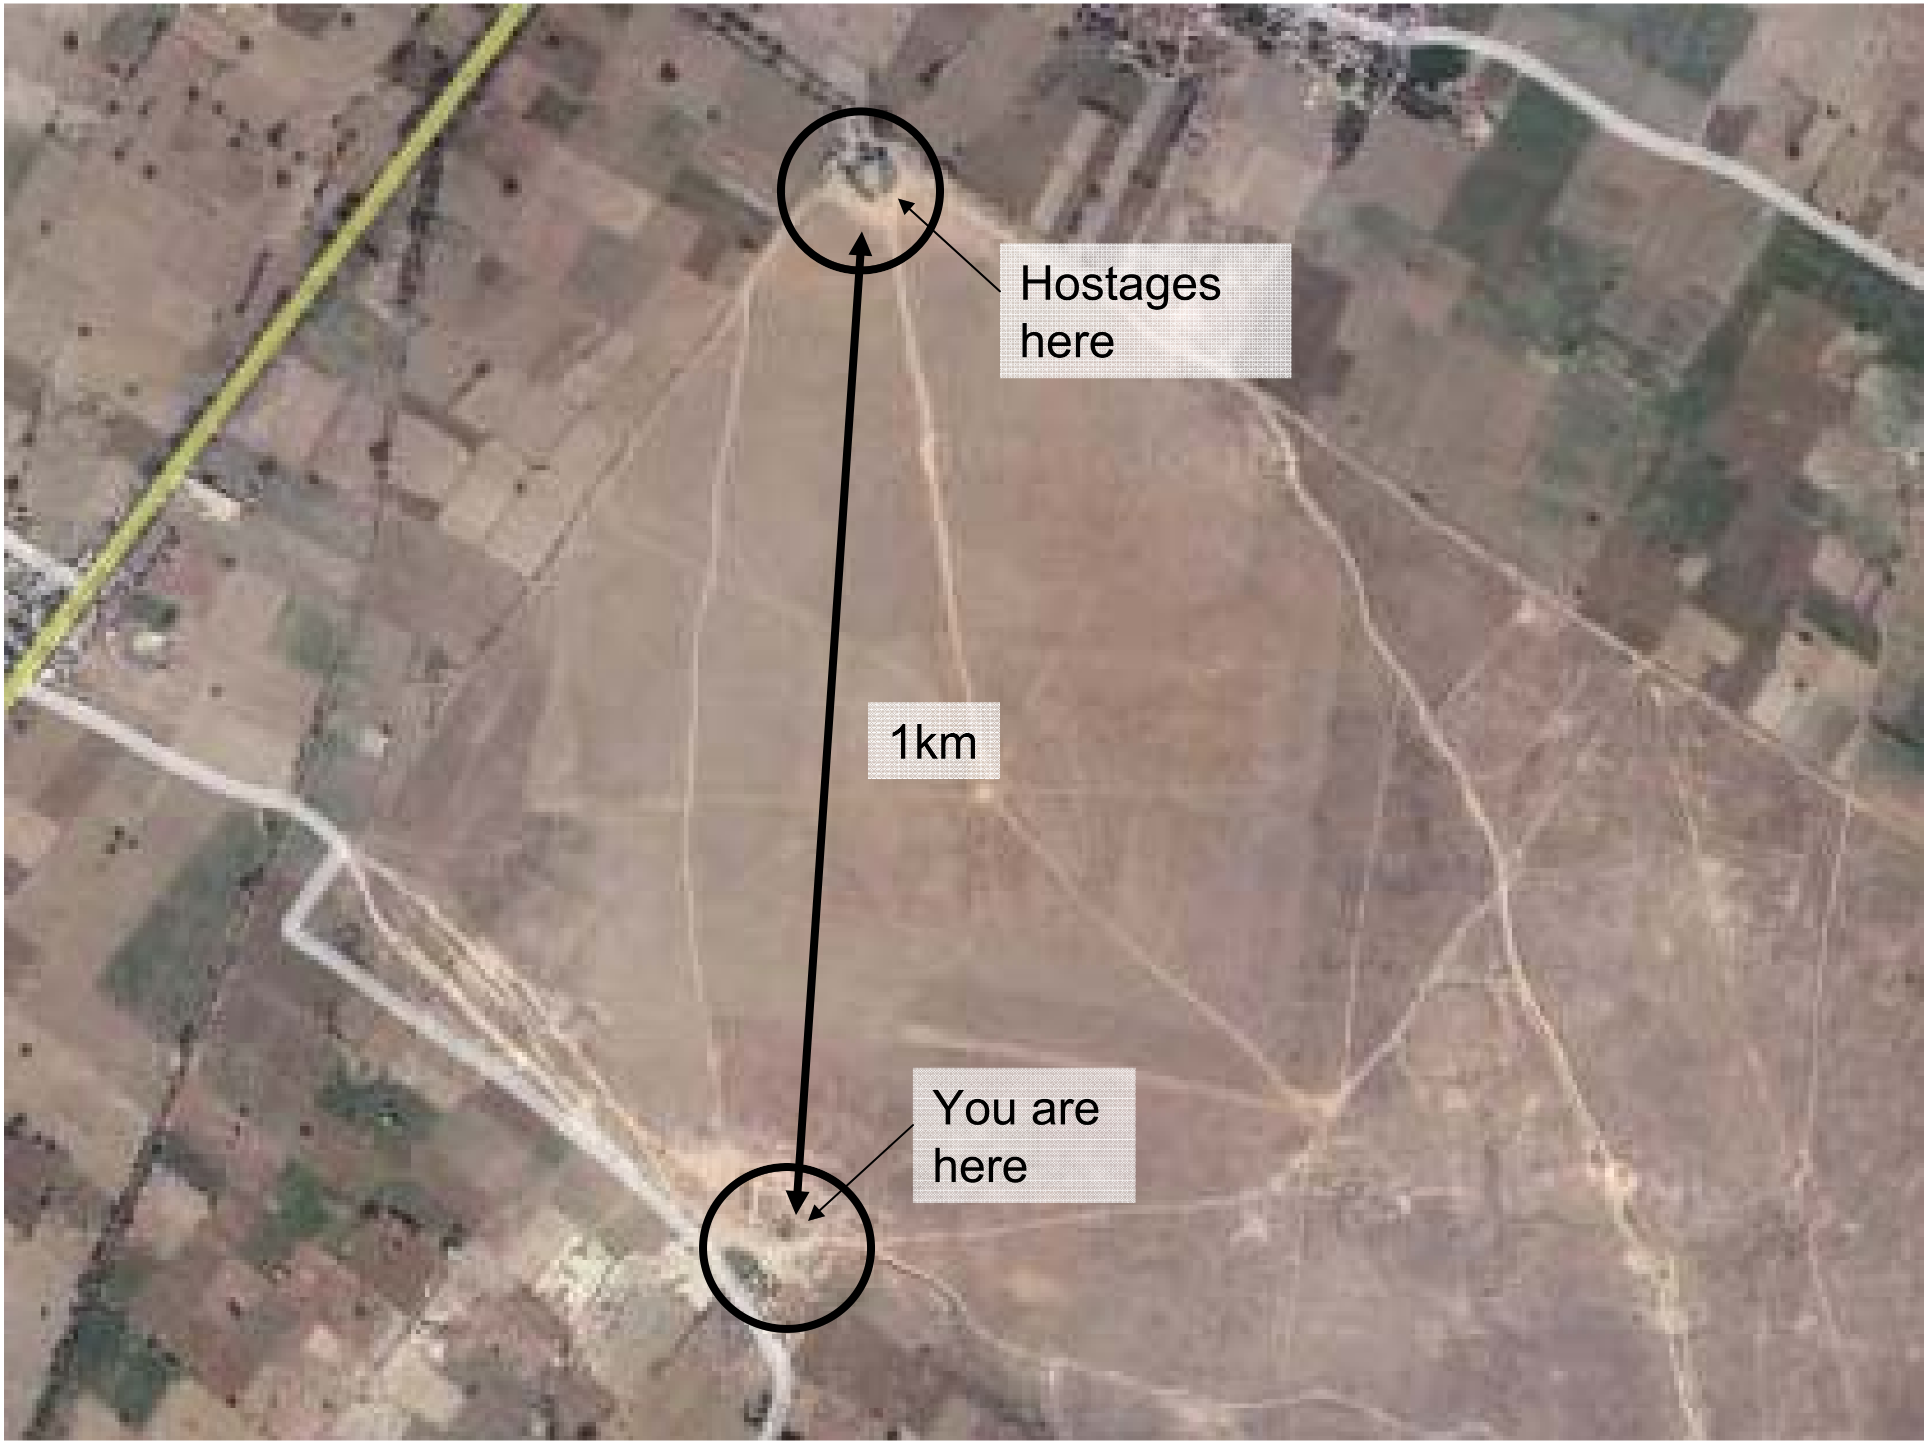
Hostages here
1km
You are here

The Mission

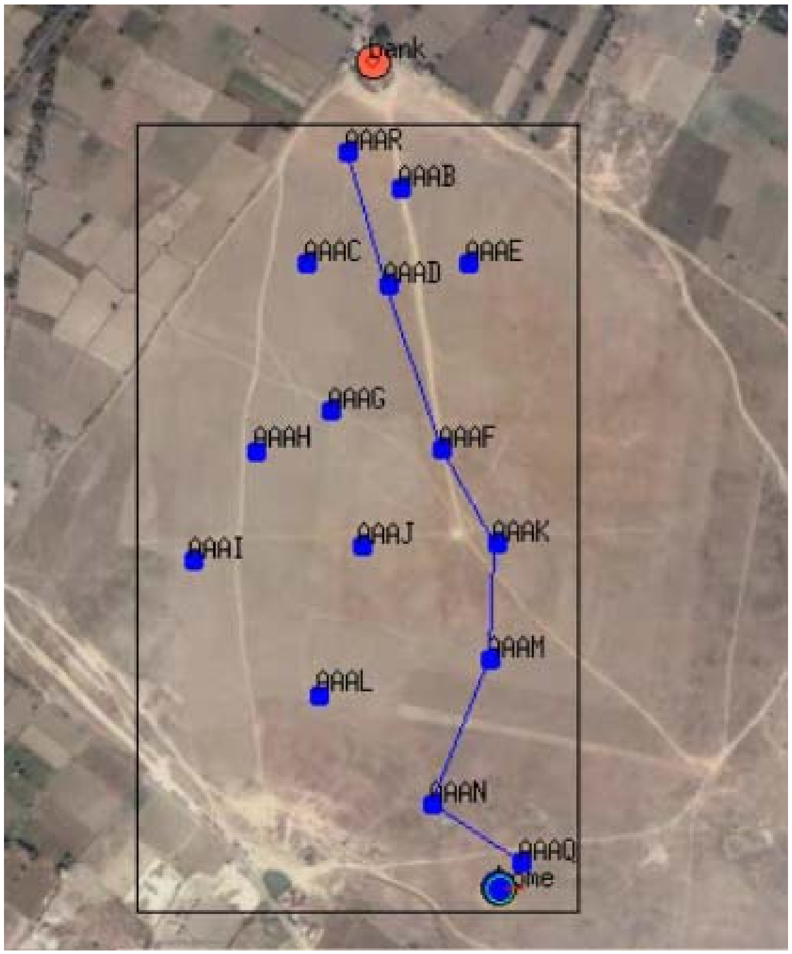
AAAR
AAAB
AAAC
AAAD
AAAE
AAAG
AAAH
AAAF
AAAI
AAAJ
AAAK
AAAM
AAAL
AAAN
AAAQ

The Mission

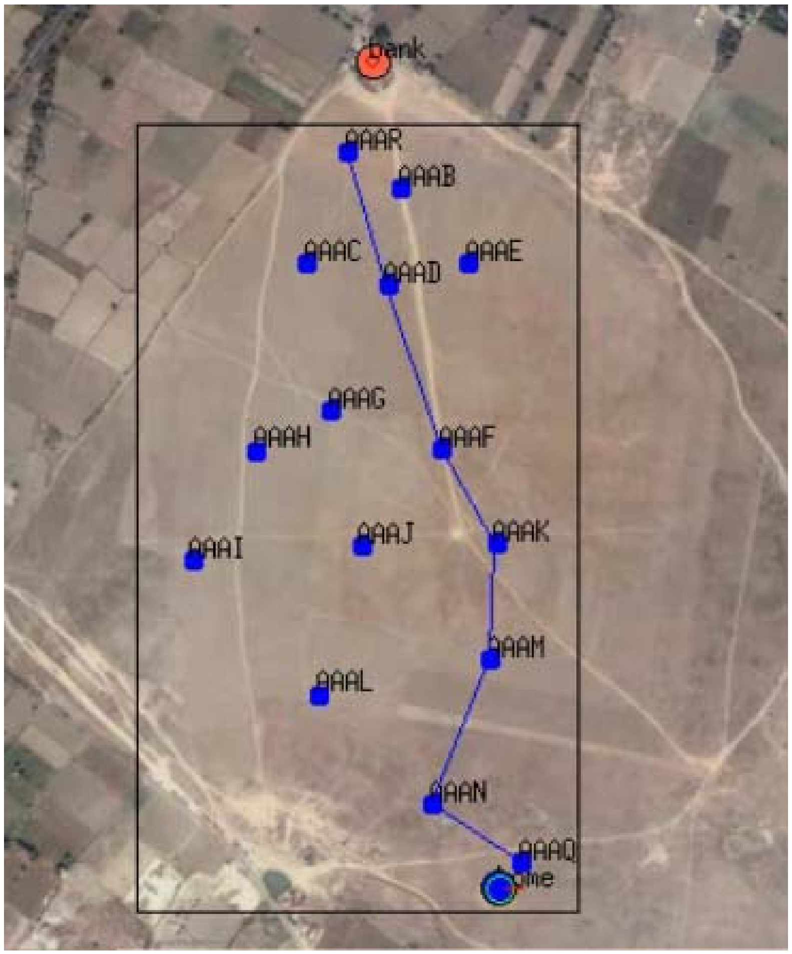
bank
AAAR
AAAB
AAAC
AAAD
AAAE
AAAG
AAAH
AAAF
AAAI
AAAJ
AAAK
AAAM
AAAL
AAAN
AAAQ

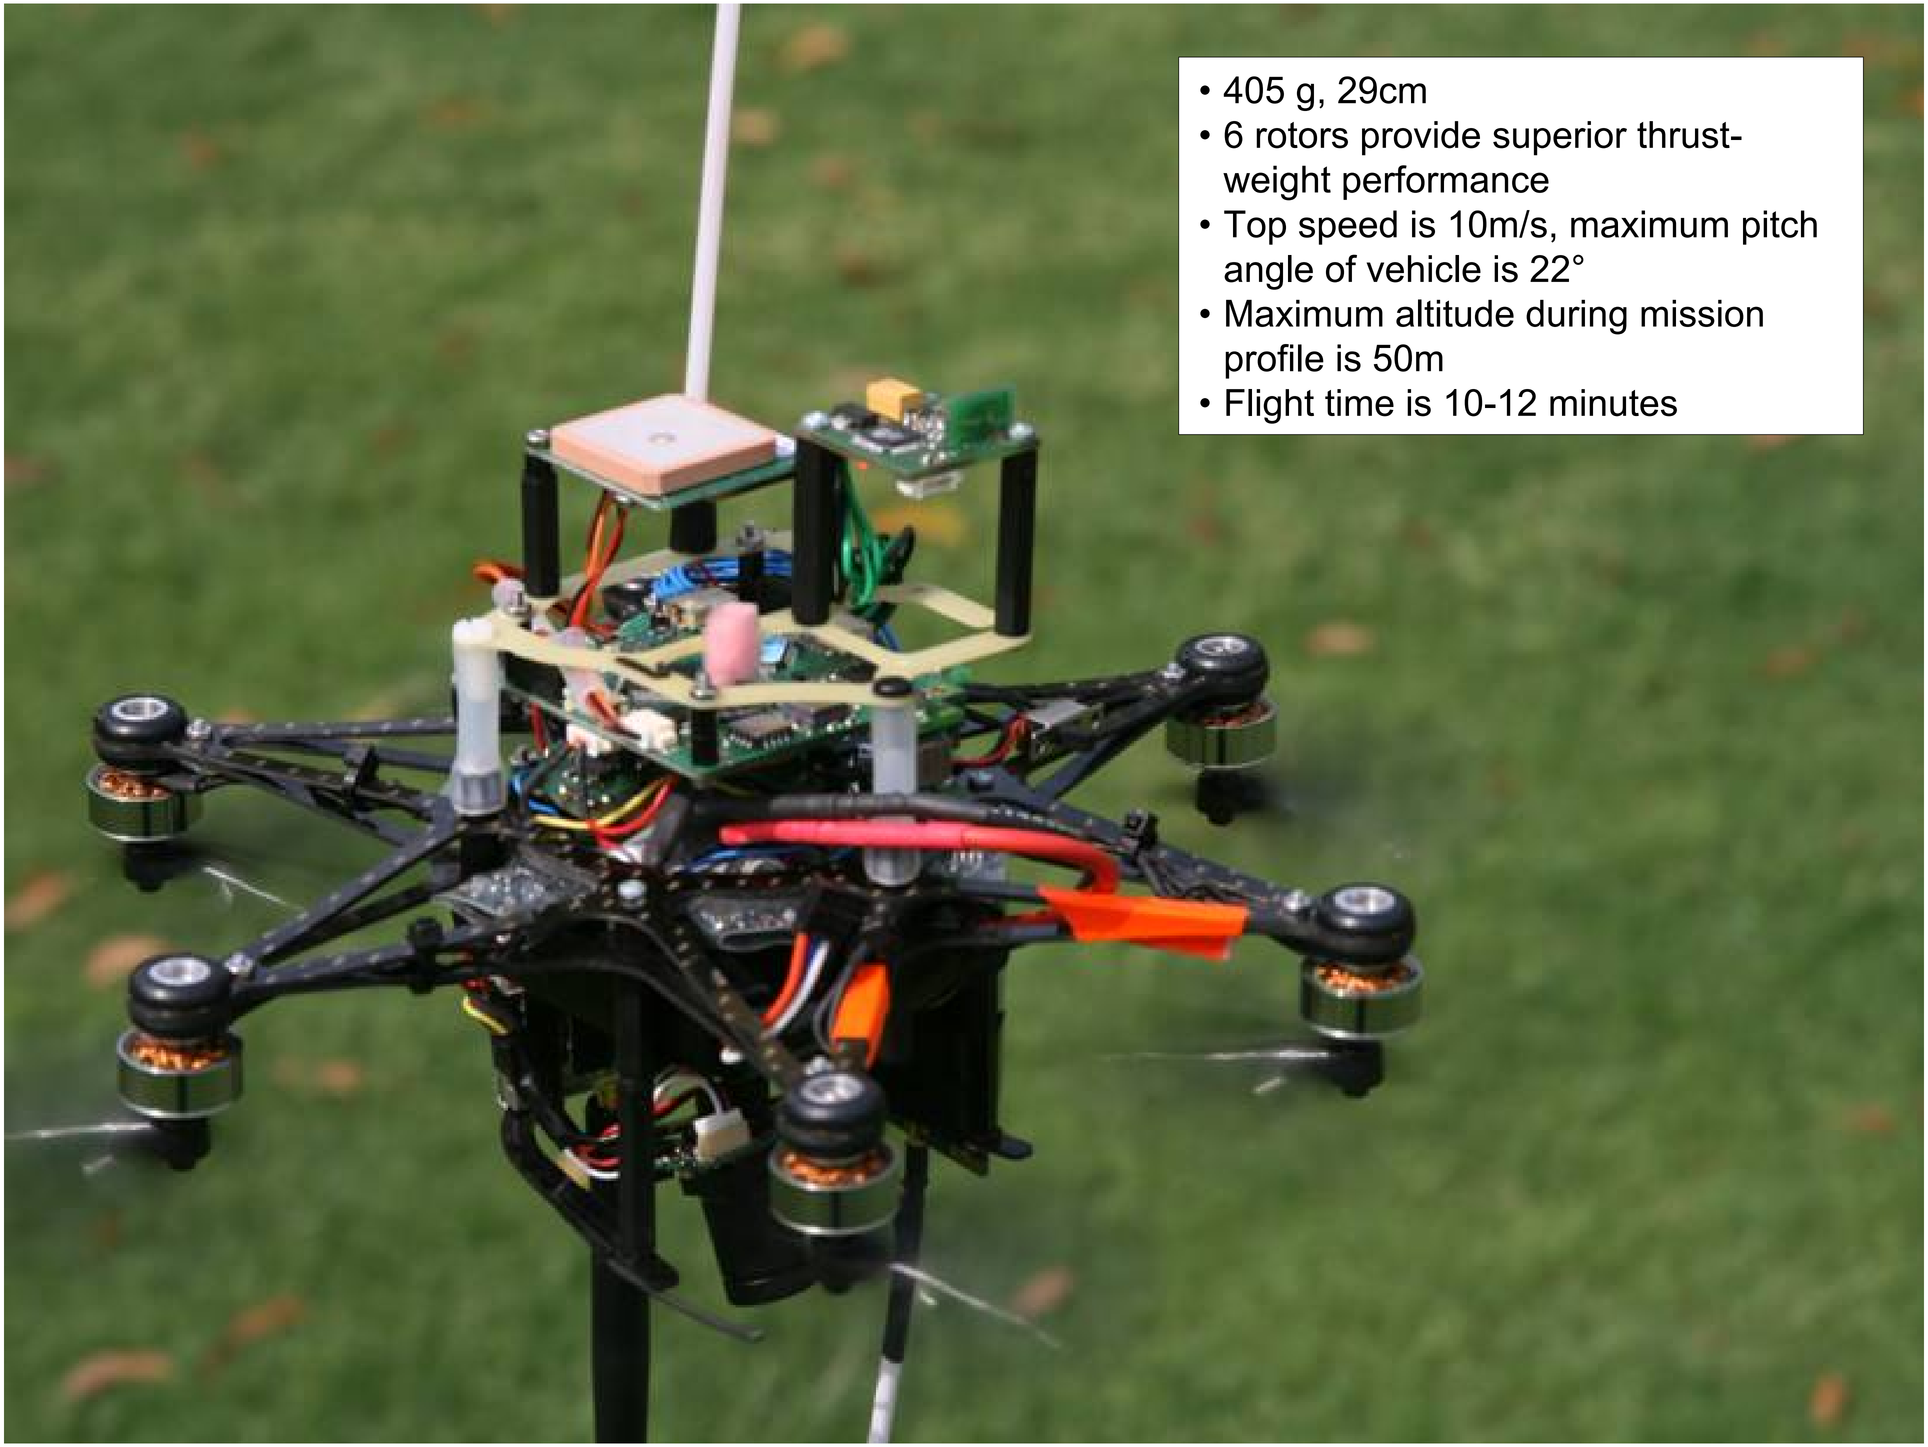

- 405 g, 29cm
- 6 rotors provide superior thrust-weight performance
- Top speed is 10m/s, maximum pitch angle of vehicle is 22°
- Maximum altitude during mission profile is 50m
- Flight time is 10-12 minutes

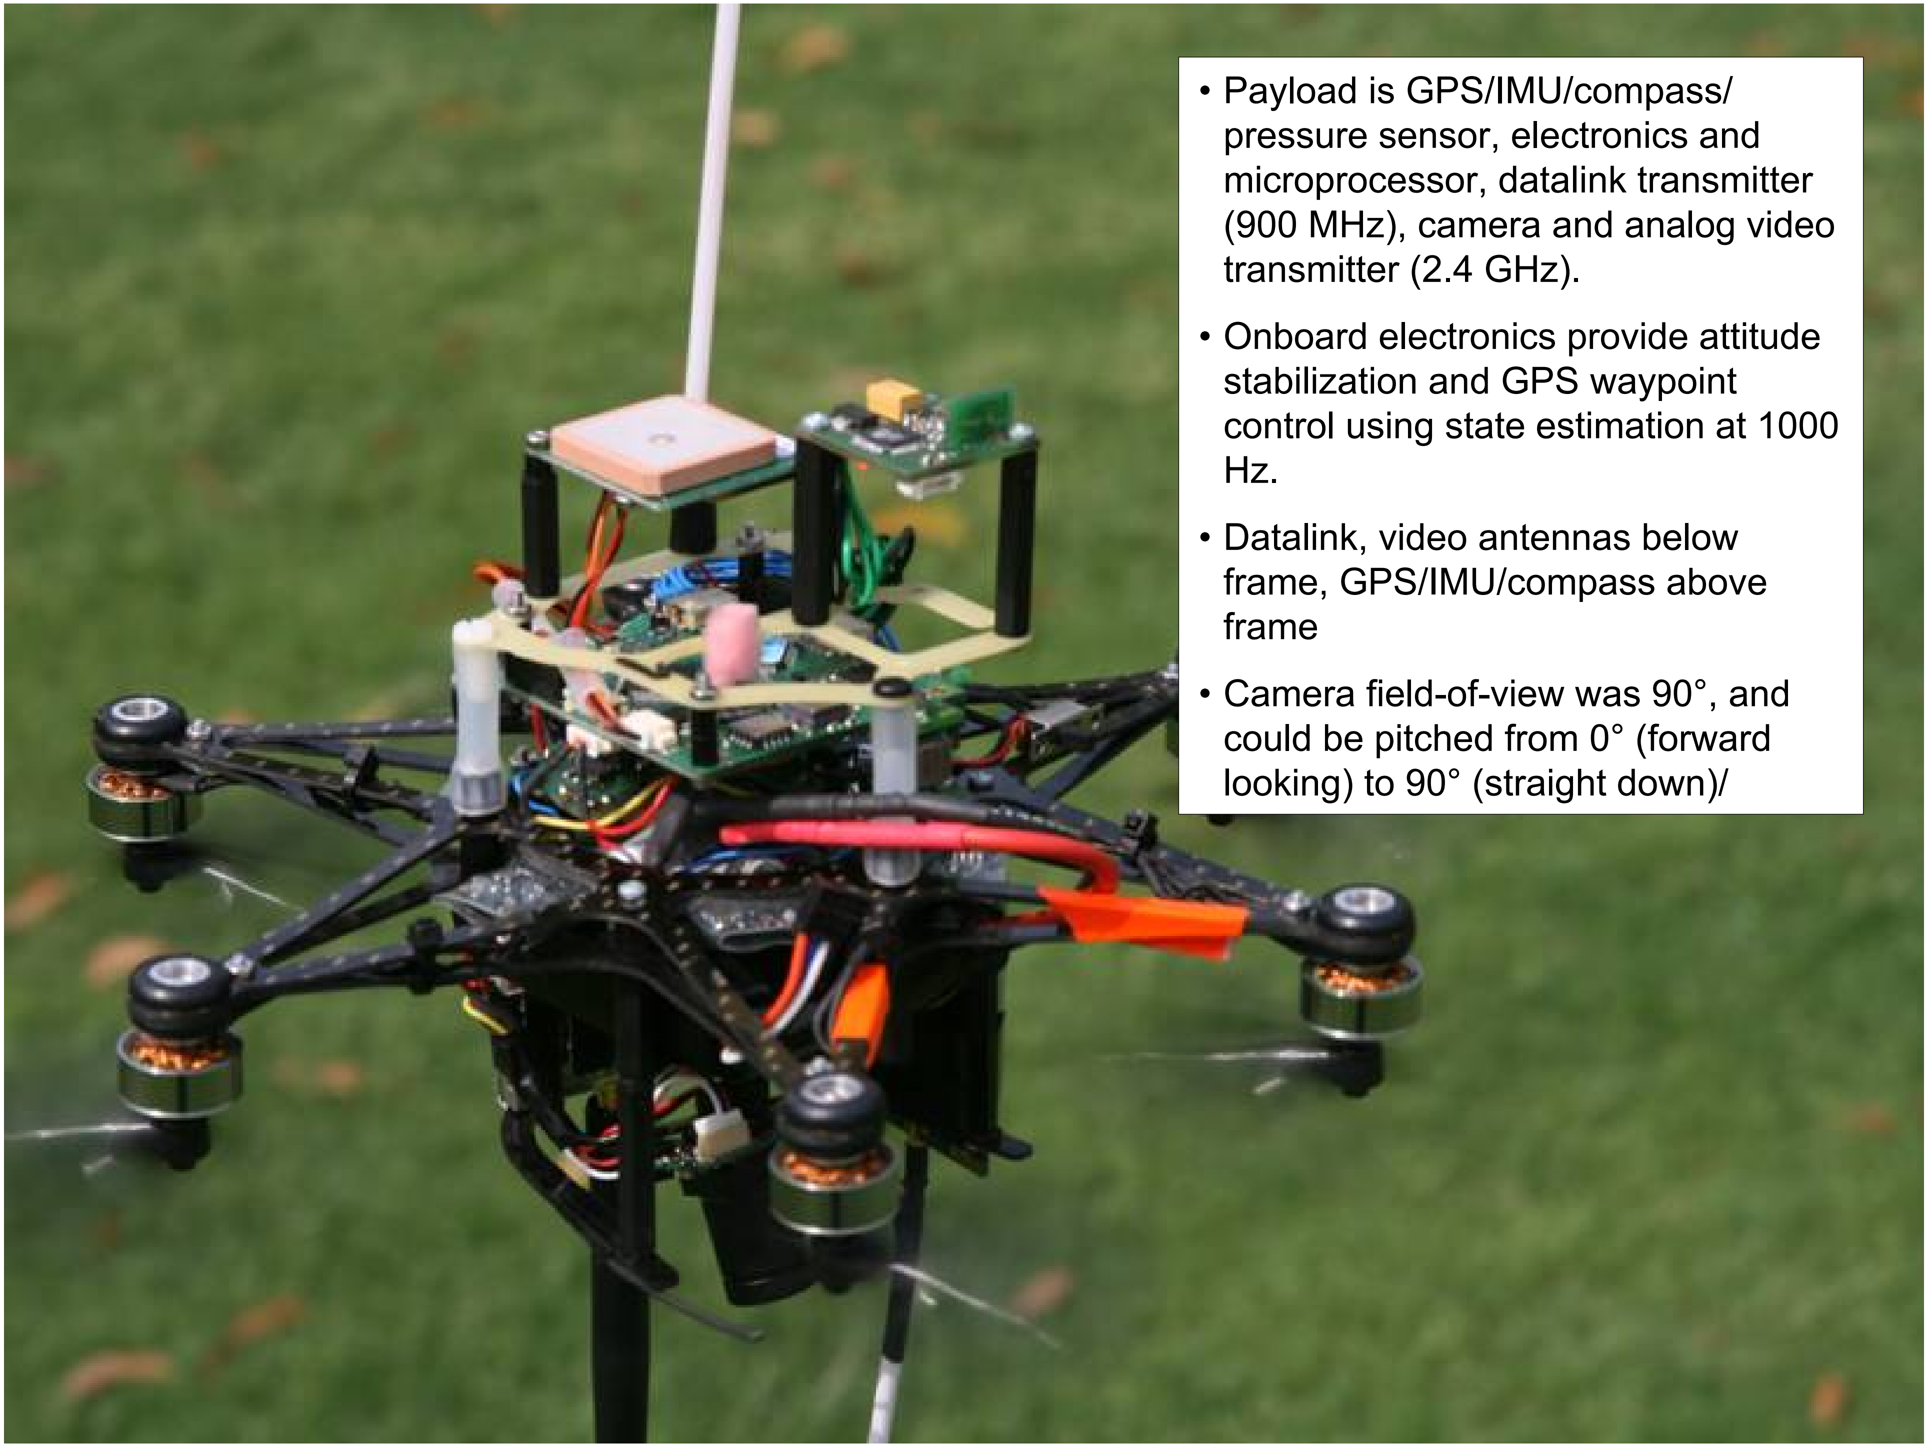

- Payload is GPS/IMU/compass/ pressure sensor, electronics and microprocessor, datalink transmitter (900 MHz), camera and analog video transmitter (2.4 GHz).
- Onboard electronics provide attitude stabilization and GPS waypoint control using state estimation at 1000 Hz.
- Datalink, video antennas below frame, GPS/IMU/compass above frame
- Camera field-of-view was 90°, and could be pitched from 0° (forward looking) to 90° (straight down)/

Attitude and Position Estimation

- Roll and pitch estimated using onboard IMU
- Yaw, X, Y Z estimated using onboard range sensors

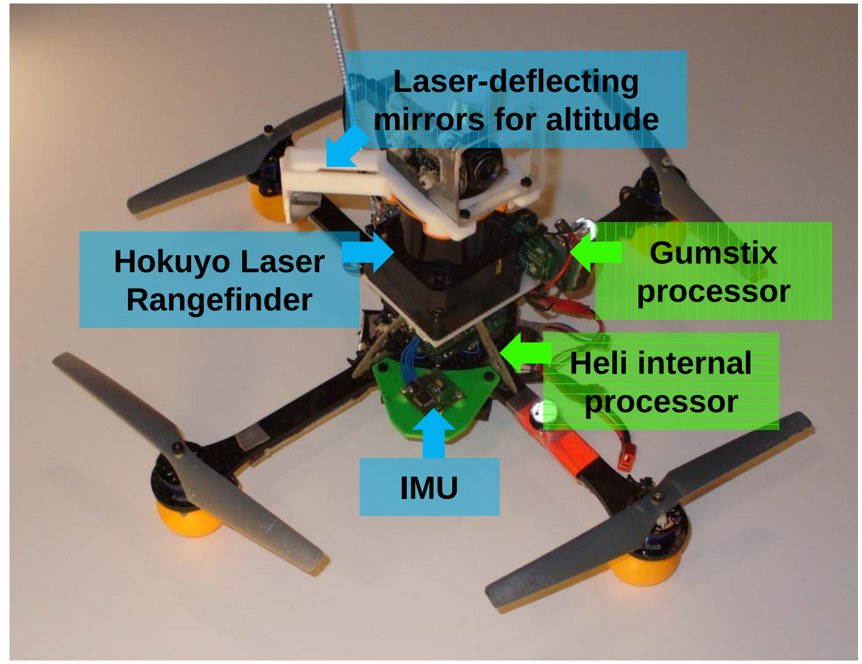

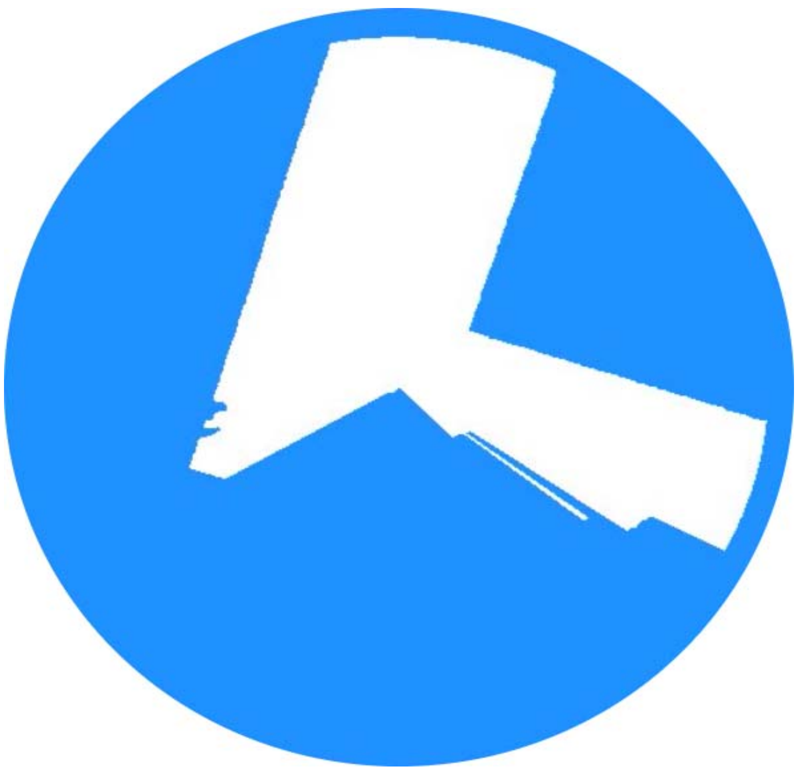

Hokuyo URG

- 4m maximum range
- 10Hz scan rate

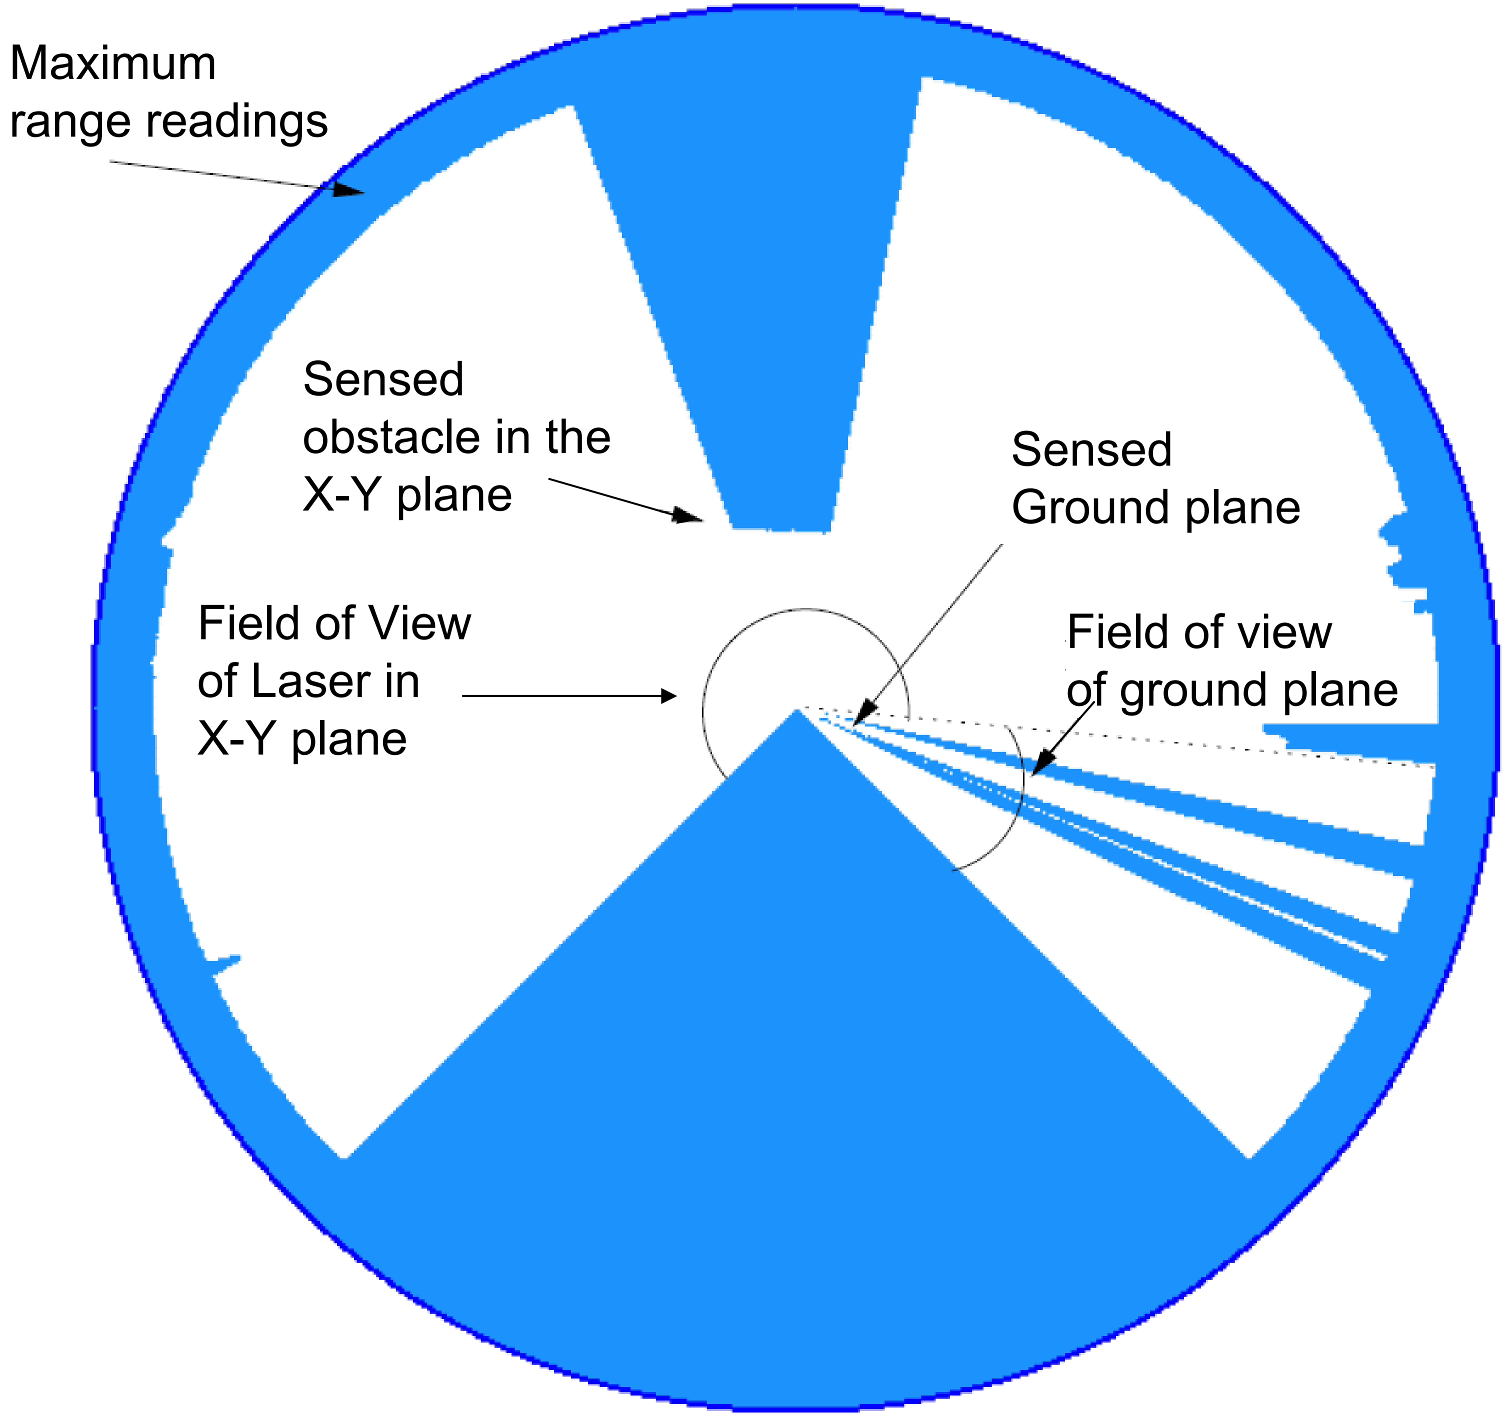
Maximum
range readings
Sensed
obstacle in the
X-Y plane
Sensed
Ground plane
Field of View
of Laser in
X-Y plane
Field of view
of ground plane

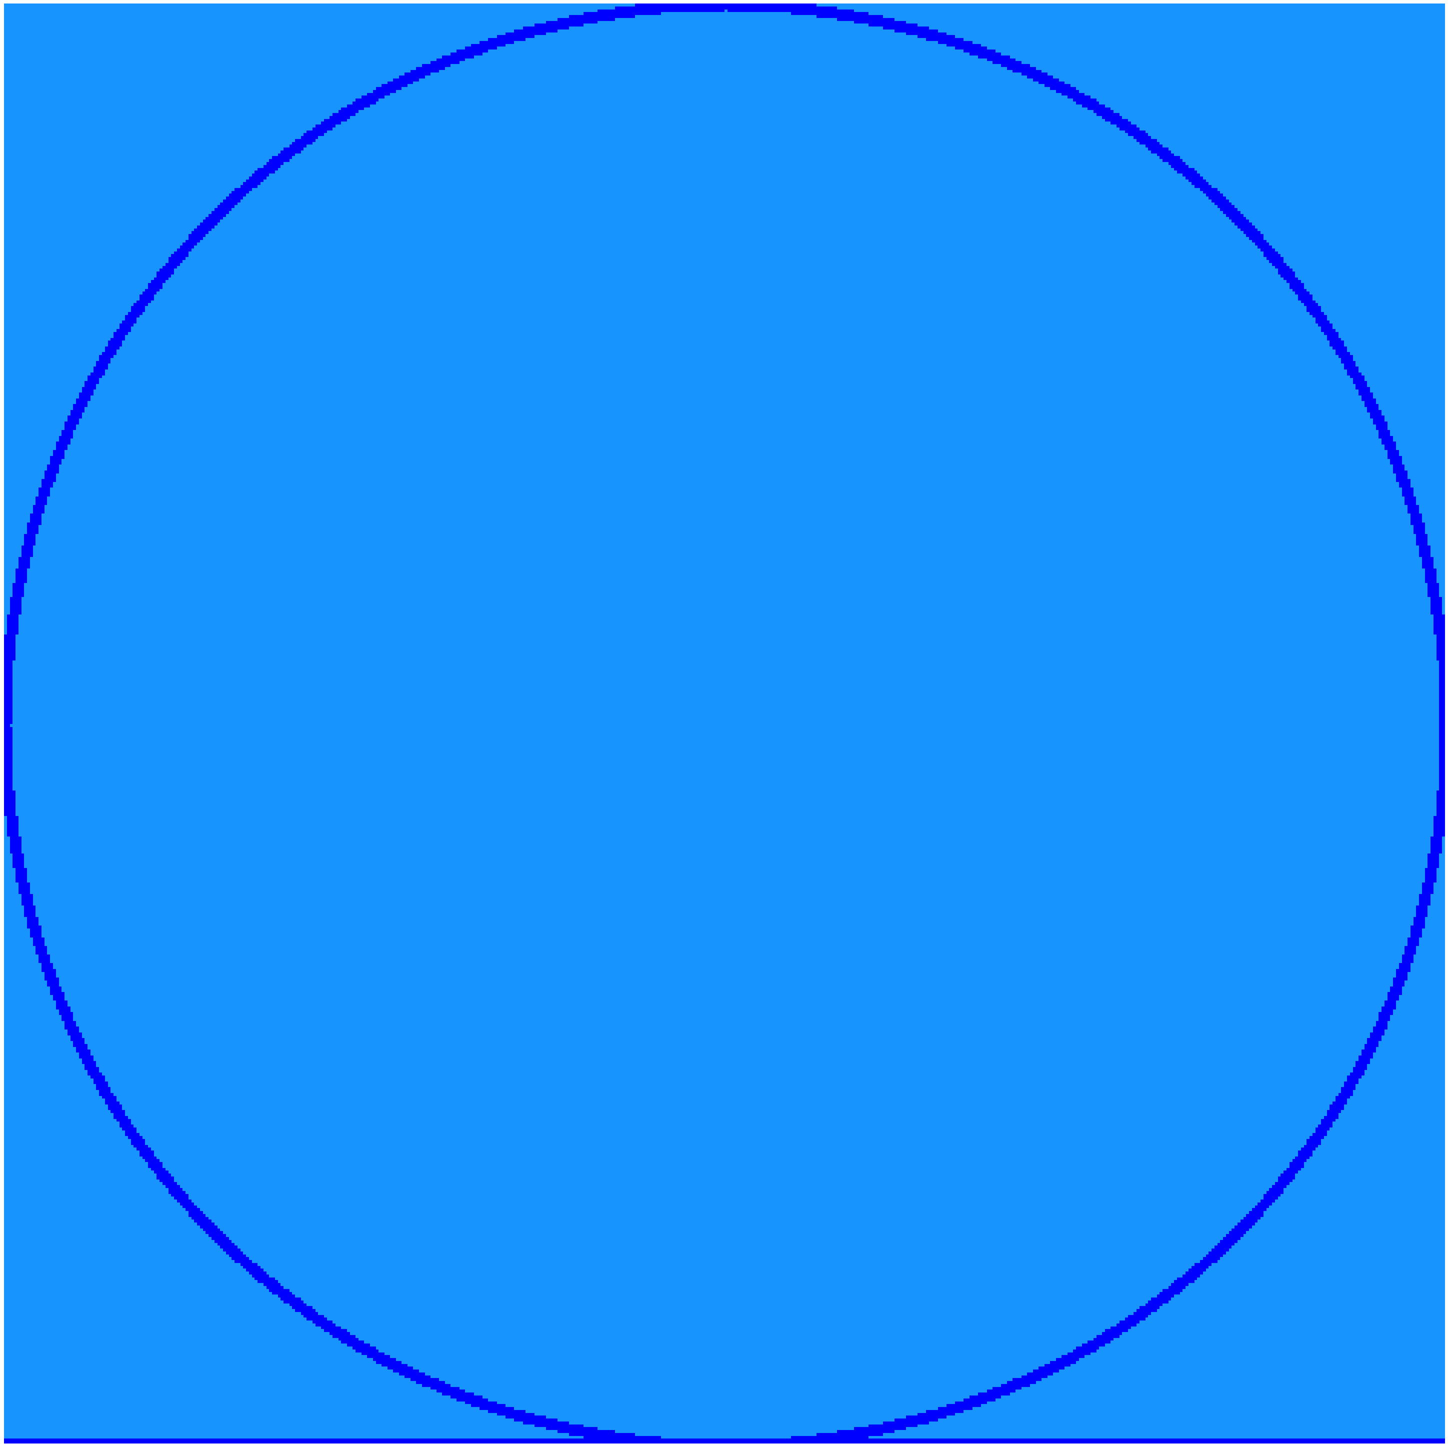

Control Models

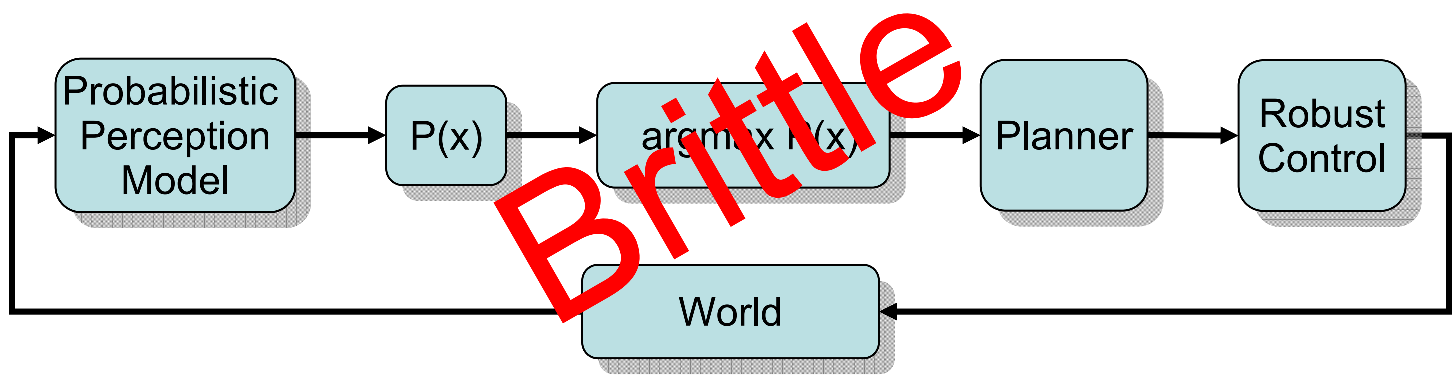

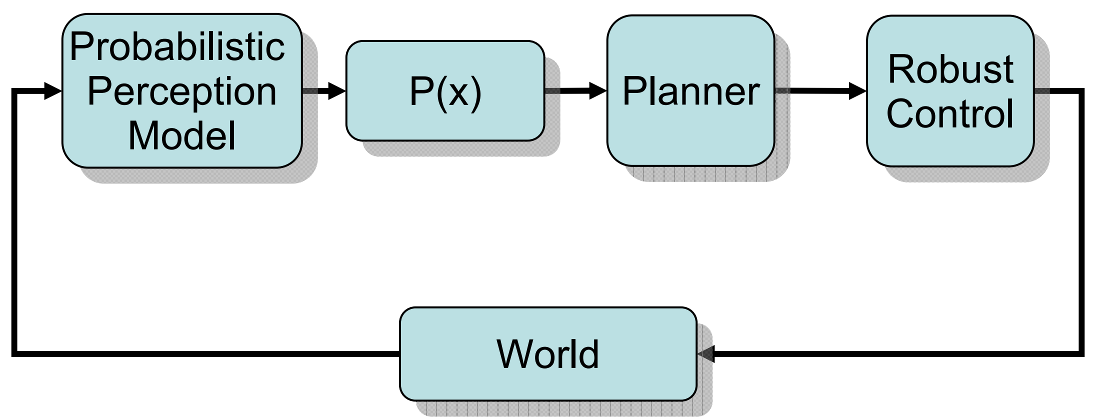

Model-uncertainty Planning

- Acting in a world in which the system has limited knowledge of the state, model of the system, or a map of the world




- Efficient inference
  - Where am I?
  - What is around me?
  - What do human team-mates want?

- Efficient planning
  - How to plan trajectories robust to sensor limitations?
  - How to explore the world?
  - How to work with human team-mates?

Sensor Limitations and Indoor Flight

- Given:
  - Map of environment

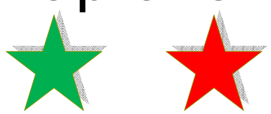

  - Start, goal locations

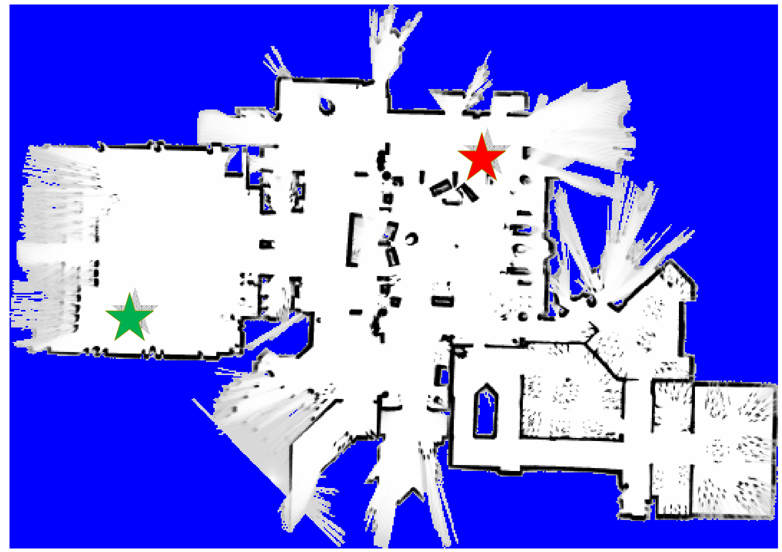

- Plan path for ***autonomous*** helicopter navigation
  - Sensor limitations

Motion Planning

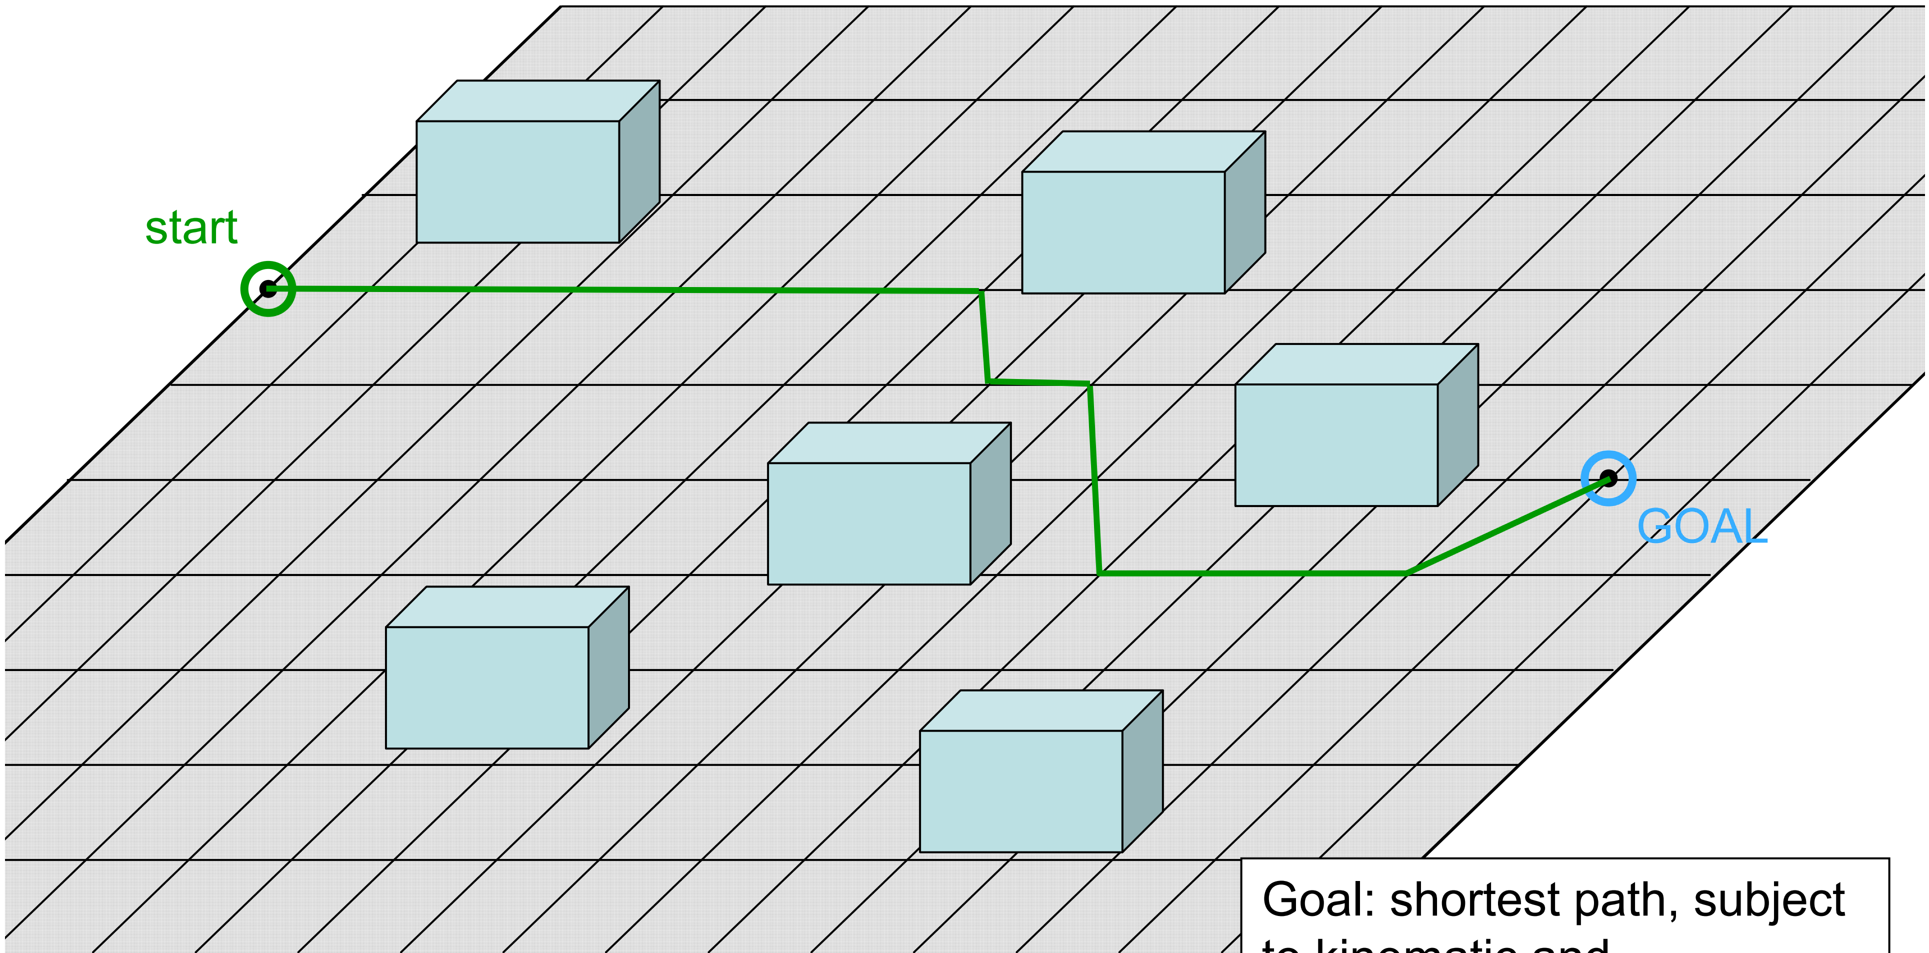

Goal: shortest path, subject to kinematic and environmental constraints

Motion Planning in High Dimensional Configuration Spaces

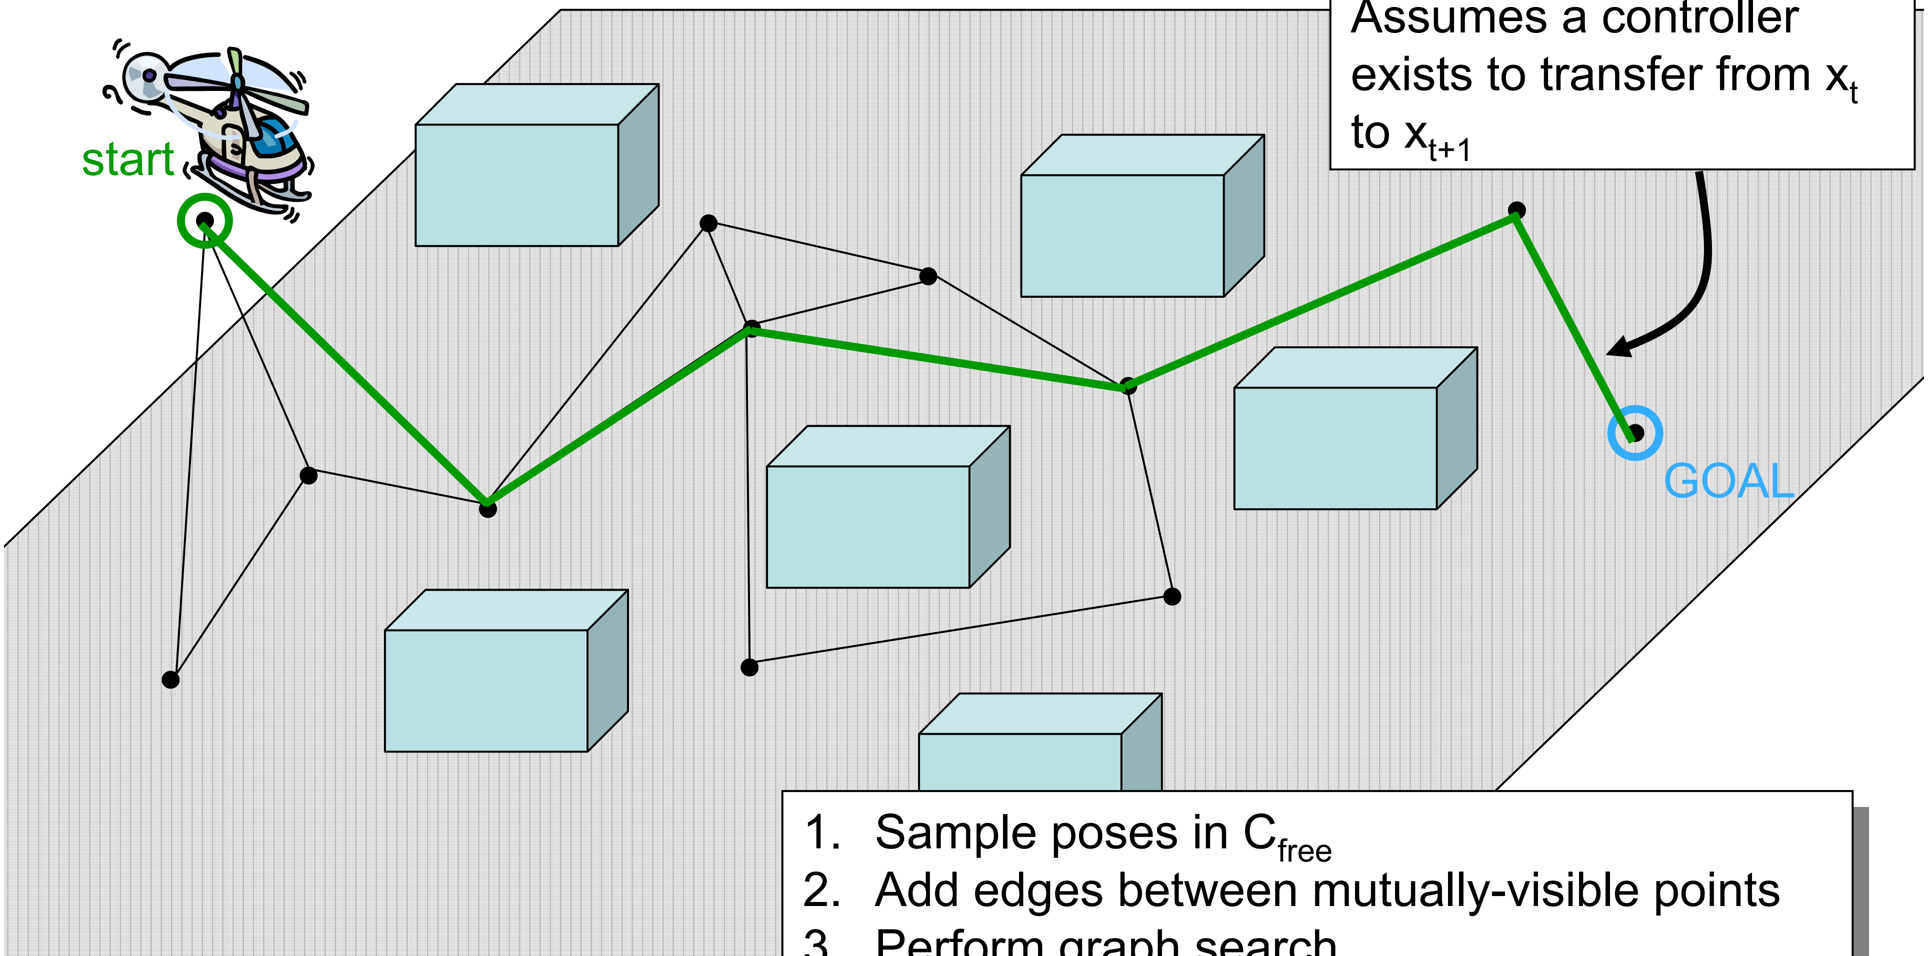

Assumes a controller exists to transfer from x_t to x_{t+1}

1. Sample poses in C_{free}
2. Add edges between mutually-visible points
3. Perform graph search

State vs. Information Space

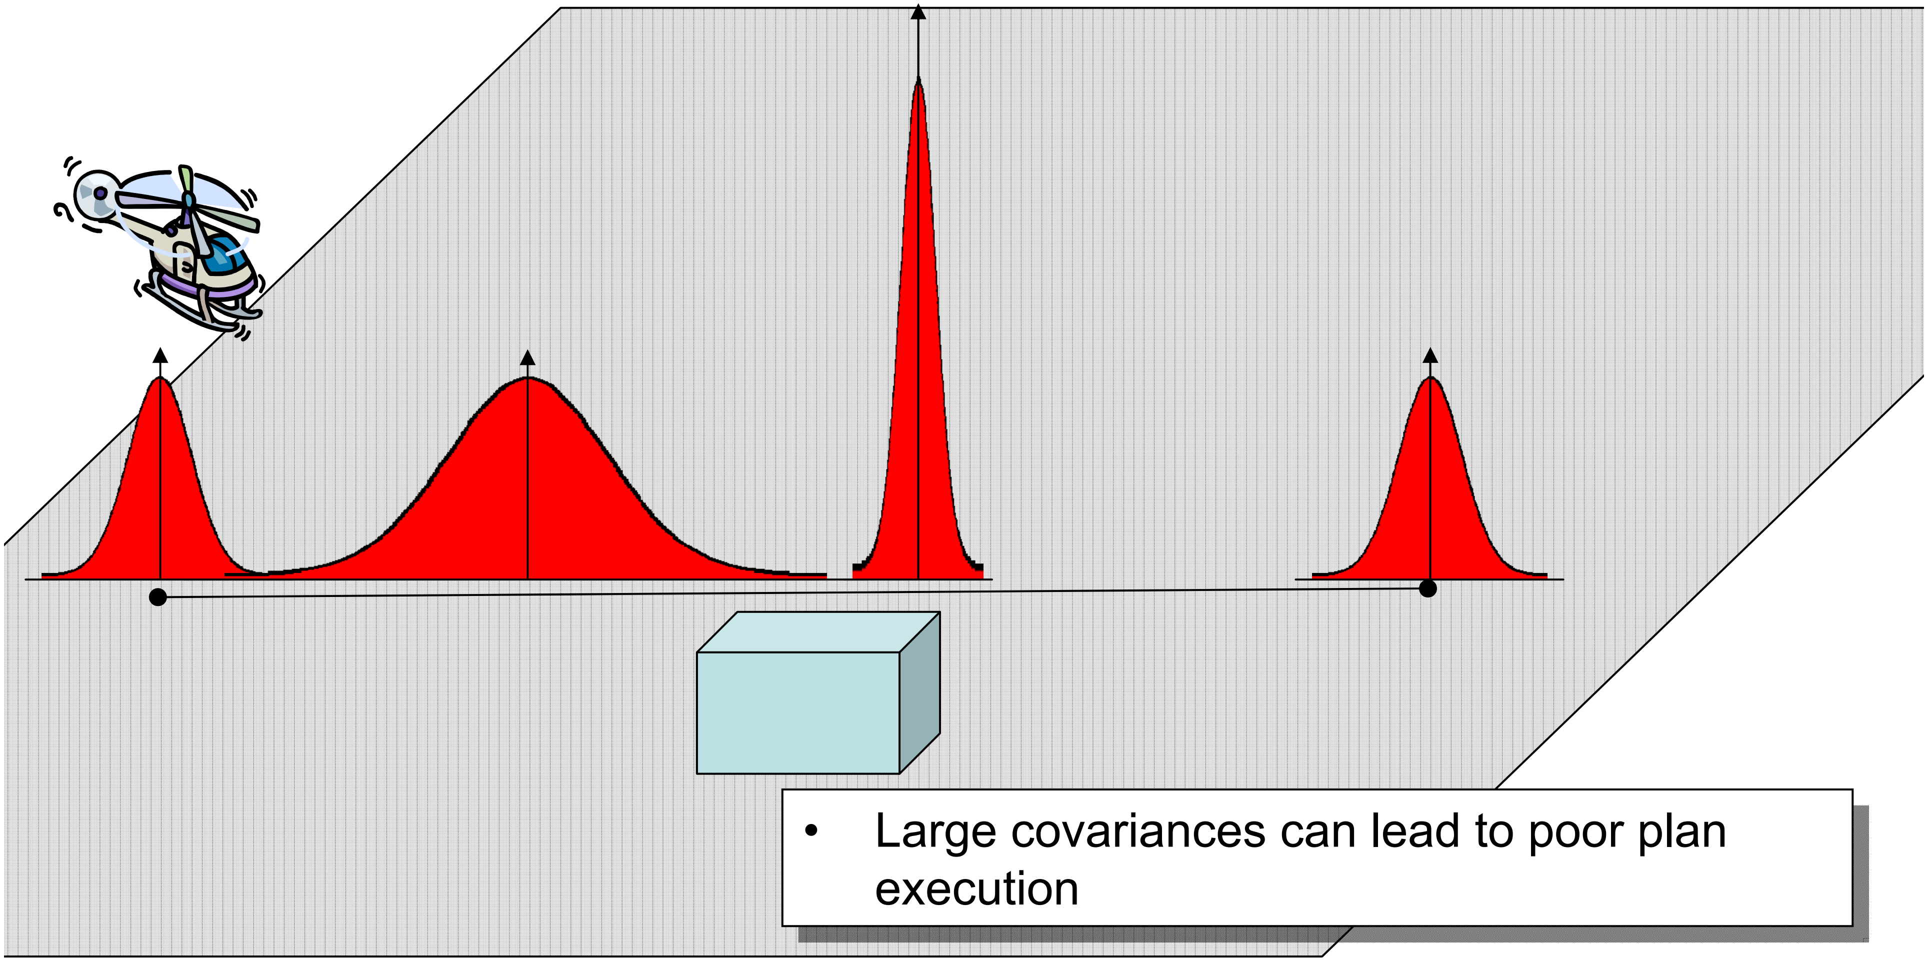

- Large covariances can lead to poor plan execution

Motion Planning in Information Space

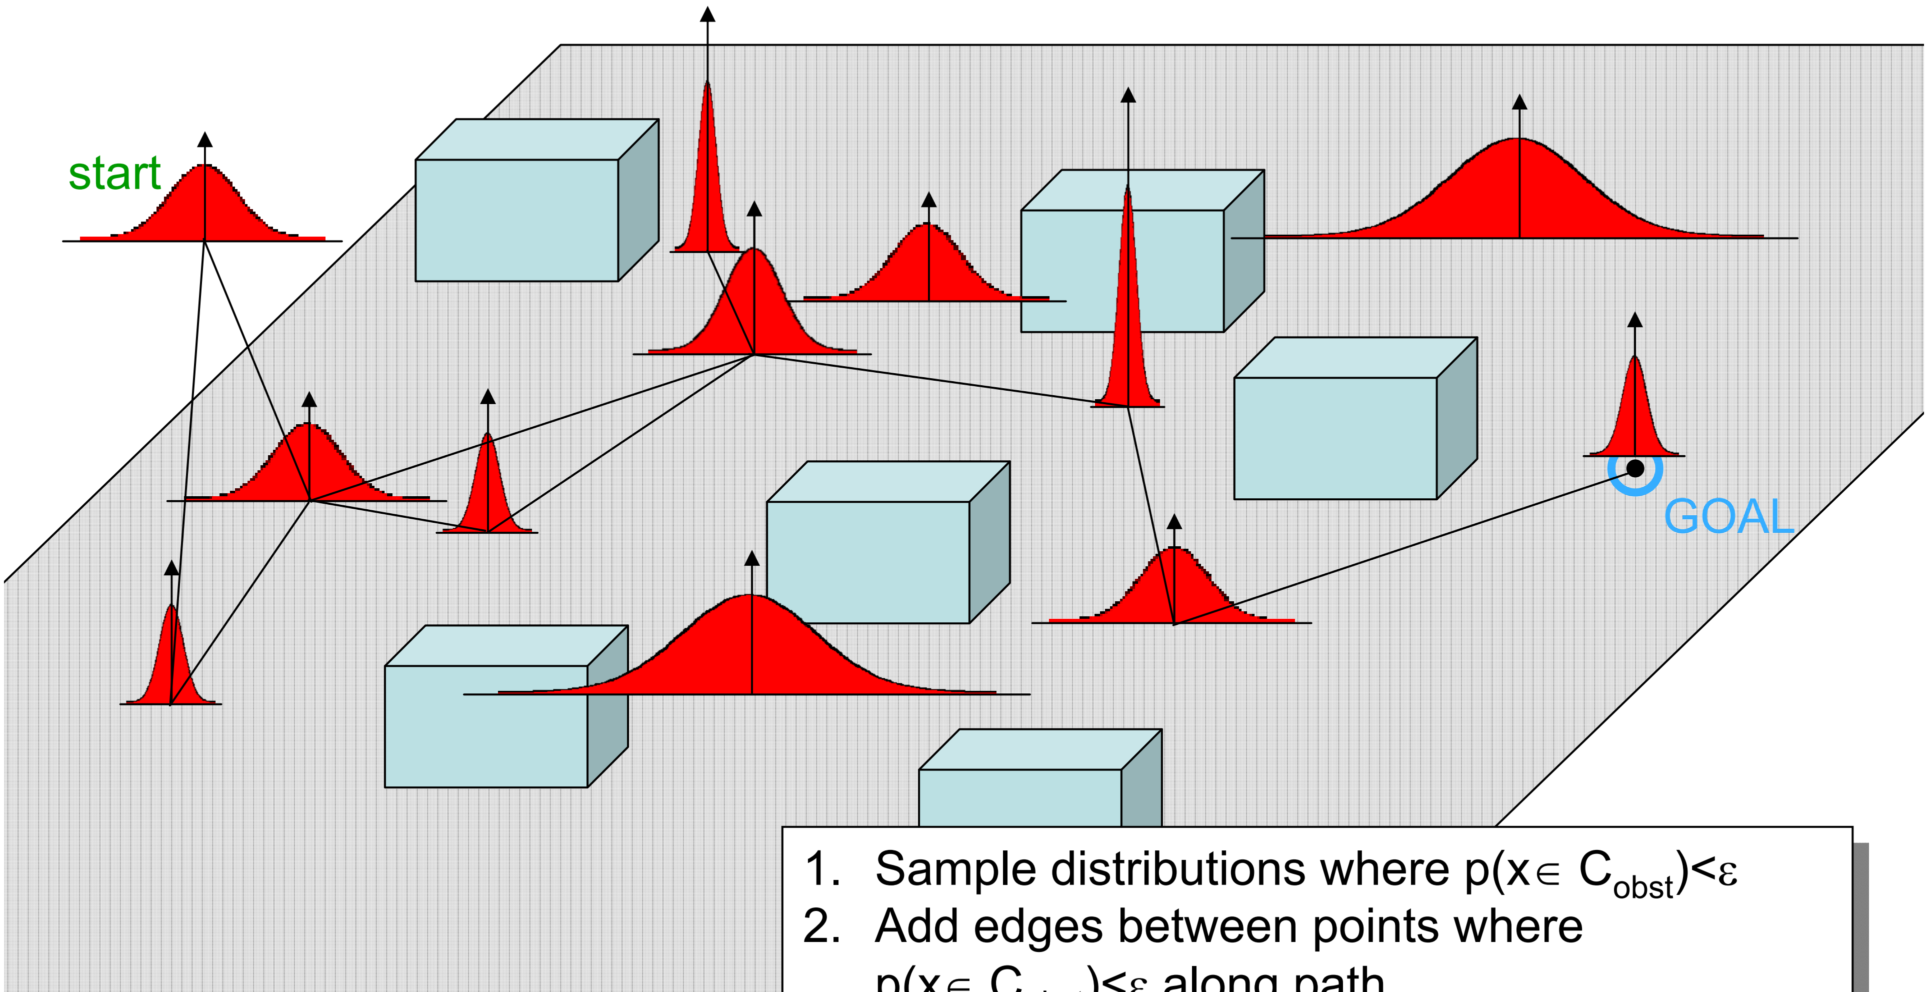

1. Sample distributions where $p(x \in C_{obst}) < \varepsilon$
2. Add edges between points where $p(x \in C_{obst}) < \varepsilon$ along path
3. Perform graph search

Problem: Edge Construction

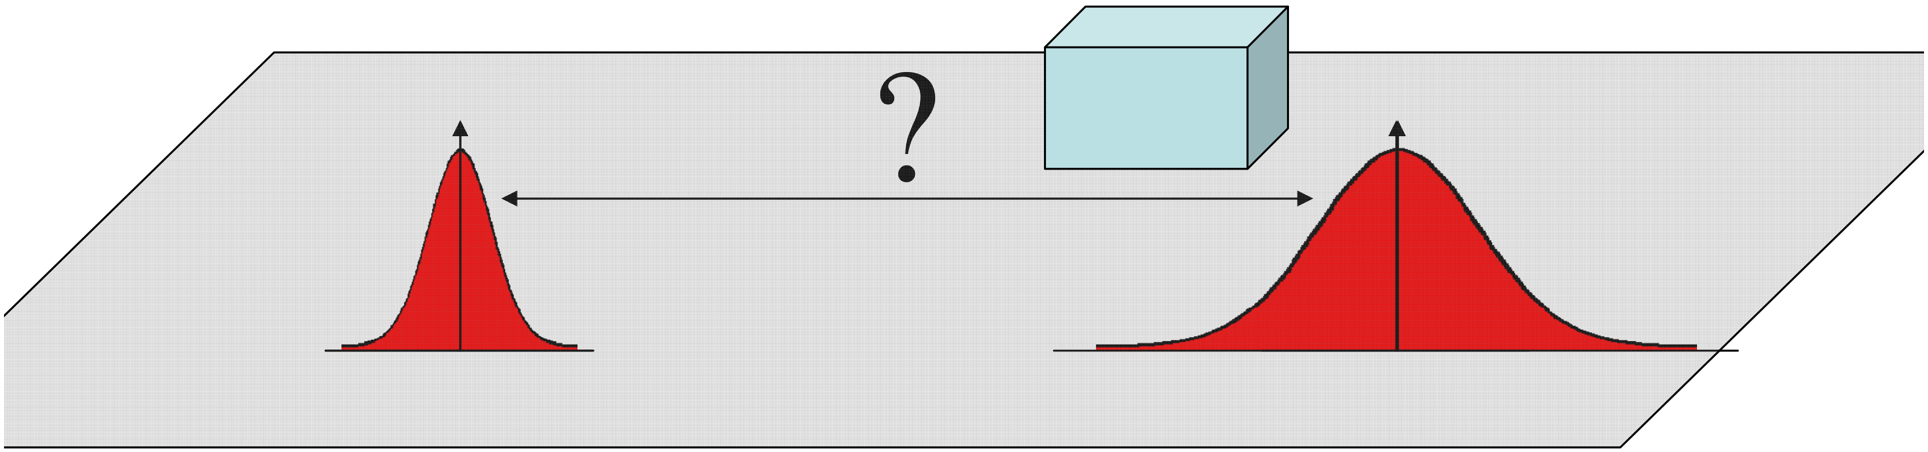

- Need $u_{0:T}$ such that $p(x|u_{0:T}) = p(x')$
- Possible solution: sample waypoints, use forward simulation to compute full posterior

Example Belief Roadmap

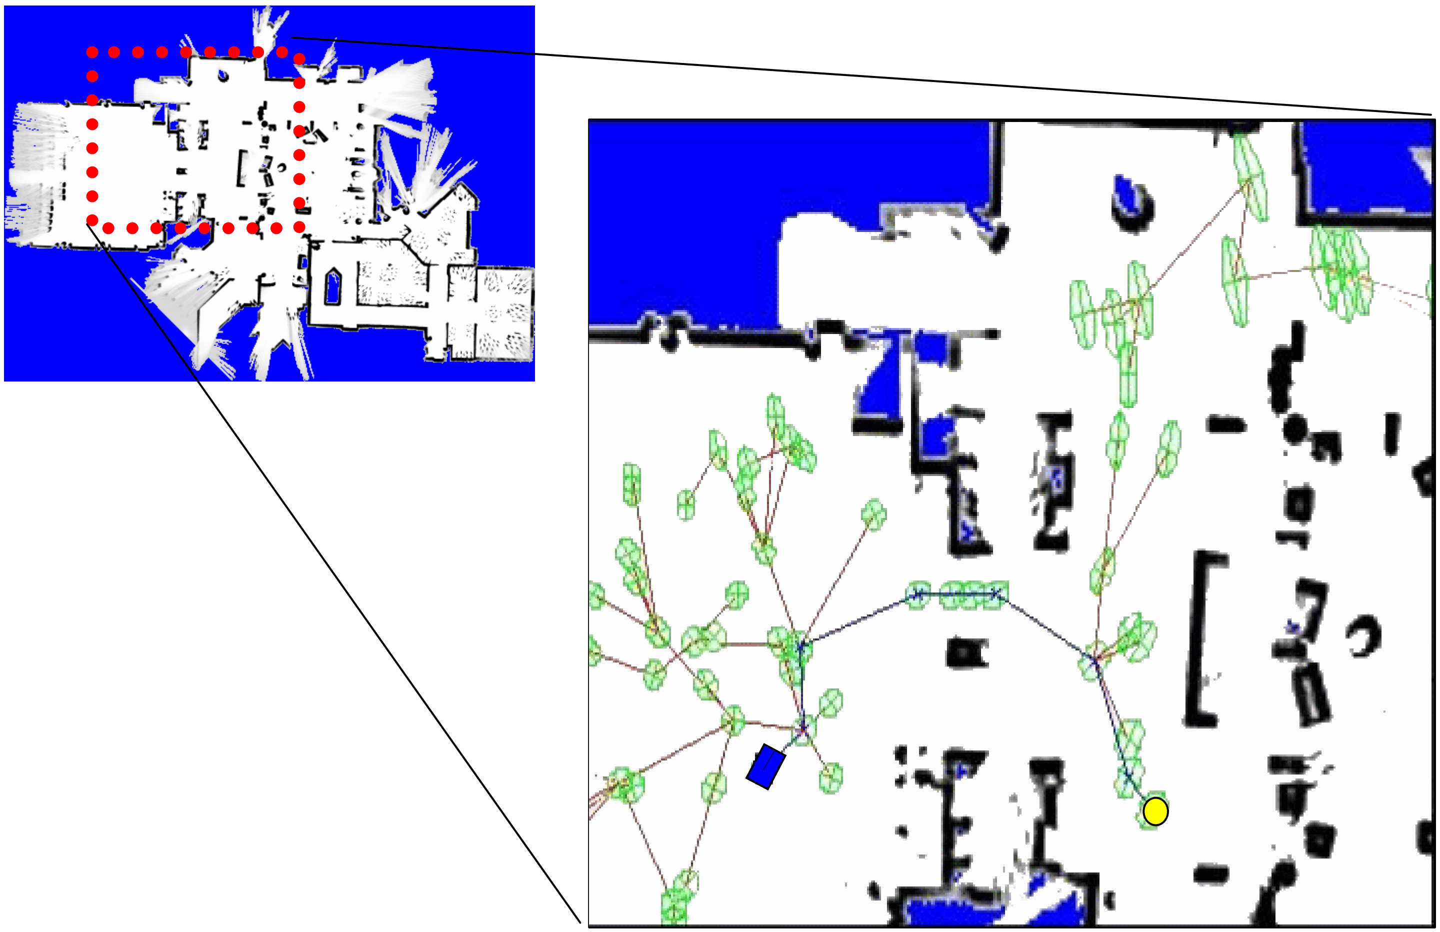

Problem: Edge Construction

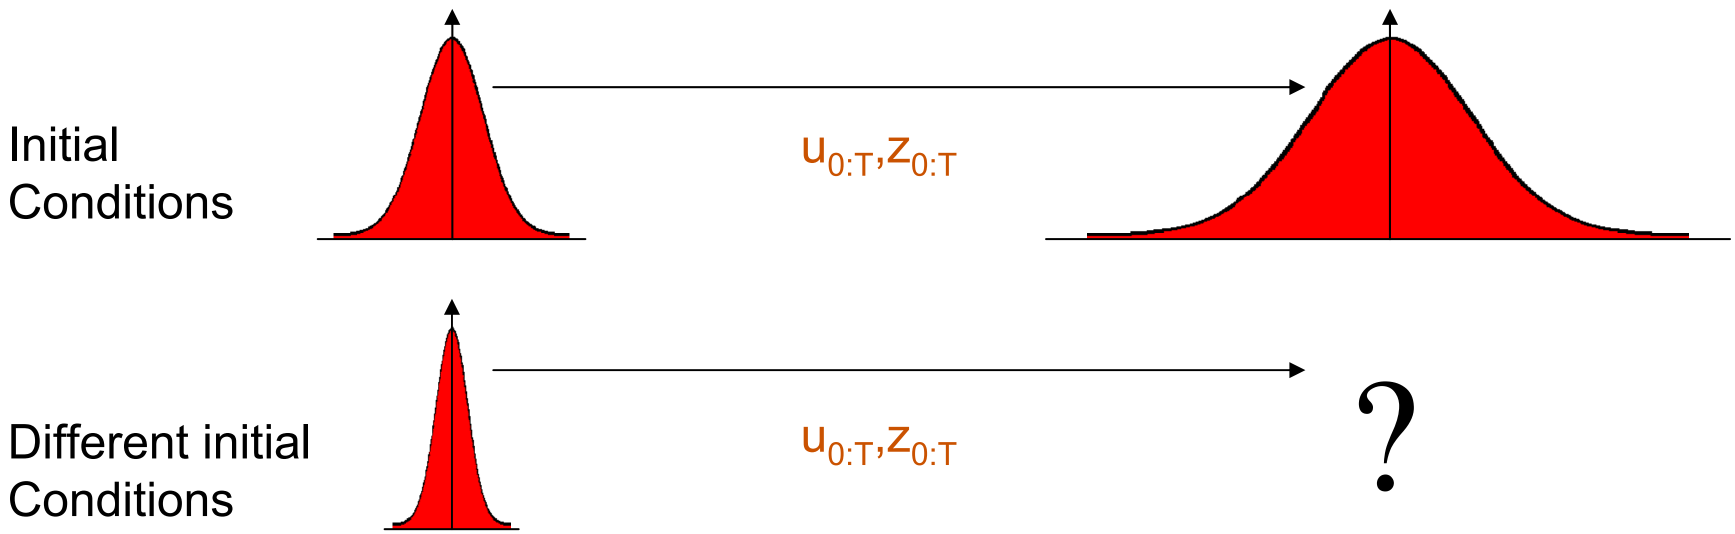

- Need to perform forward simulation (and belief prediction) along each edge for every start state
- Computing minimum cost path of 30 edges: ≈100 seconds
- Not an issue for single queries: clearly a problem for multi-query planning

Multi-Step Update as One-Step

EKF Covariance Update

Control: $\overline{\Sigma}_t = G\Sigma_{t-1}G^T + R$

Measurement: $\Sigma_t = \left(\overline{\Sigma}_t^{\,-1} + HQ^{-1}H^T\right)^{-1}$

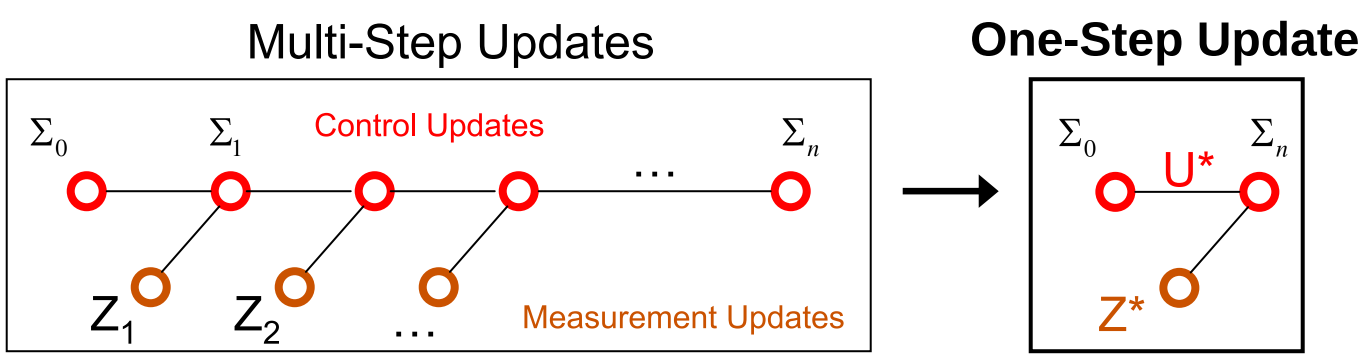

Solution: Decomposition

- Key idea: factor the covariance matrix

$$\Sigma = BC^{-1}$$

- Motion update

$$\overline{\Sigma}_t = \overline{E}_t \overline{D}_t^{-1}$$

$$\begin{bmatrix} \overline{D}_t \\ \overline{E}_t \end{bmatrix} = \begin{bmatrix} 0 & G_t^{-T} \\ G_t & R_t G_t^{-T} \end{bmatrix} \begin{bmatrix} B_{t-1} \\ C_{t-1} \end{bmatrix}$$

Solution: Decomposition

- Key idea: factor the covariance matrix

$$\Sigma = BC^{-1}$$

- Measurement update

$$\Sigma_t = B_t C_t^{-1}$$

$$\begin{bmatrix} B_t \\ C_t \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 1 & M_t \end{bmatrix} \begin{bmatrix} \overline{D}_t \\ \overline{E}_t \end{bmatrix}$$

Solution: Decomposition

- One-step transfer function for the covariance:

$$\zeta_t = \begin{bmatrix} 0 & I \\ I & M_t \end{bmatrix} \begin{bmatrix} 0 & G_t^{-T} \\ G_t & R_t G_t^{-T} \end{bmatrix}$$

$$\Rightarrow \begin{bmatrix} B_T \\ C_T \end{bmatrix} = \left(\prod_{t=0}^{T} \zeta_t \right) \begin{bmatrix} B_0 \\ C_0 \end{bmatrix}$$

- (To recover covariance, $\Sigma = BC^{-1}$)
- This trick is not new.
  - Kaileth et al., Linear State Estimation.
  - Mourikis and Roumeliotis, 2006.

The Belief Roadmap Algorithm

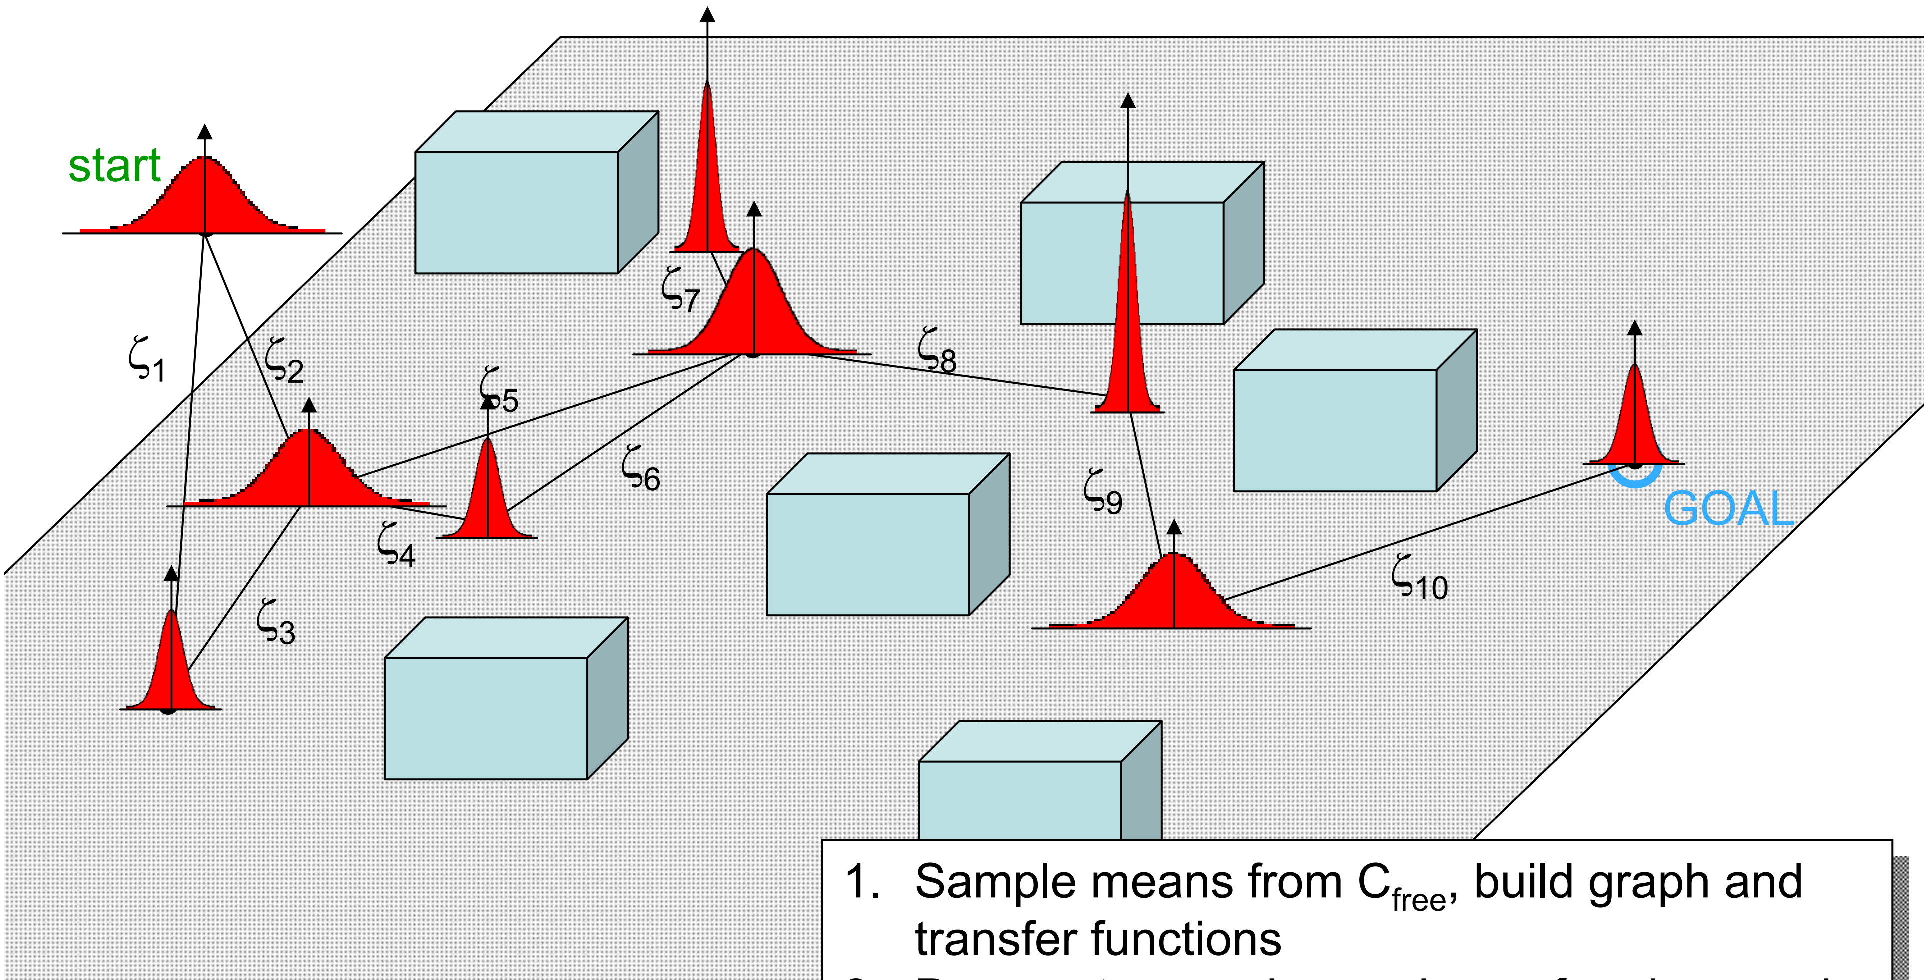

1. Sample means from C_{free}, build graph and transfer functions
2. Propagate covariances by performing graph search

The Belief Roadmap

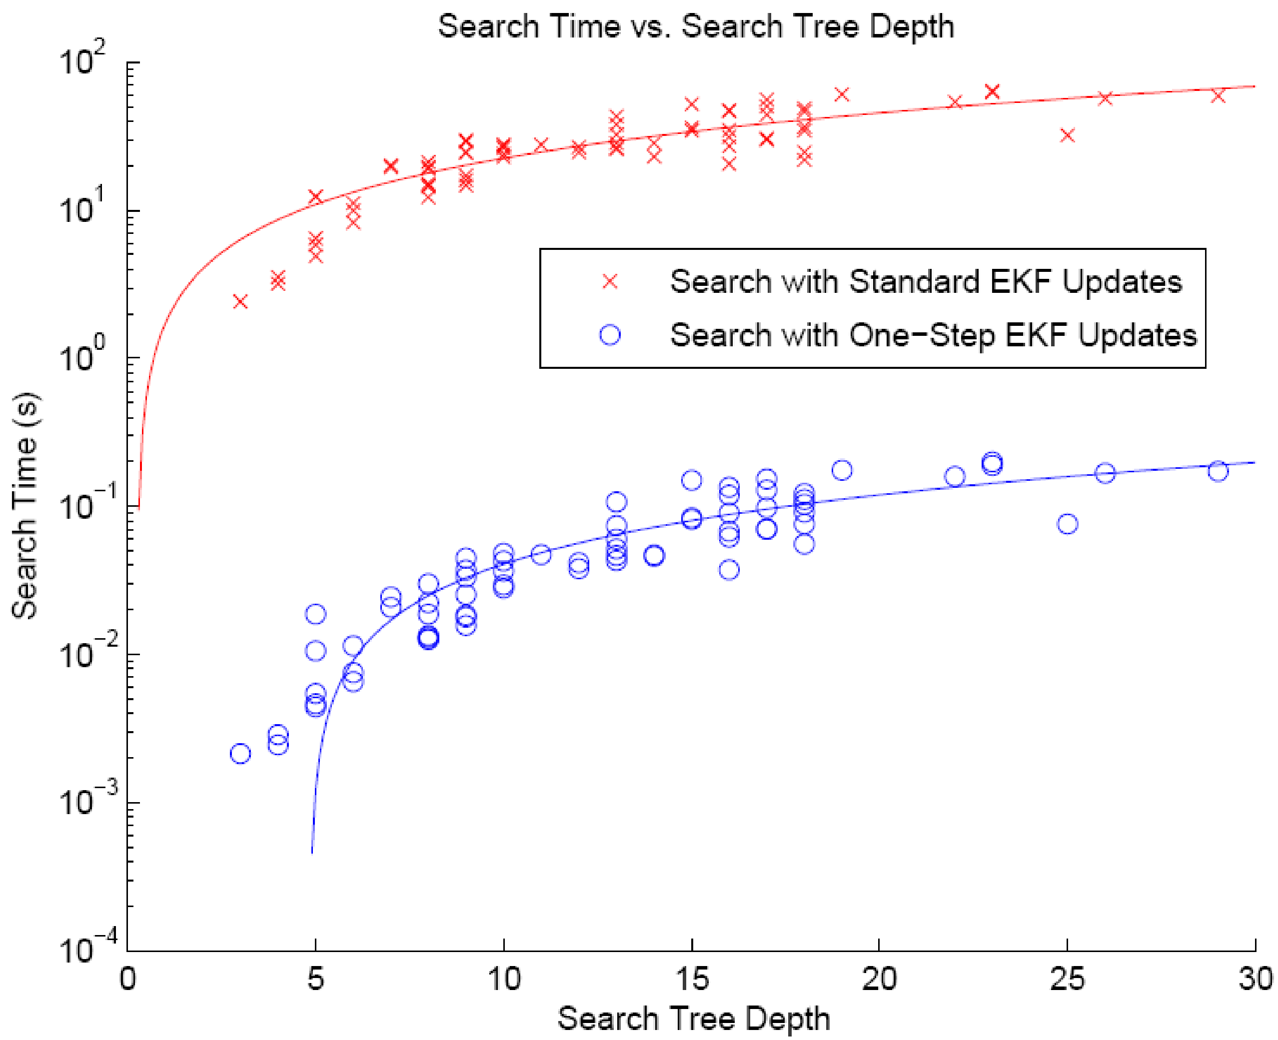

Improving Sampling

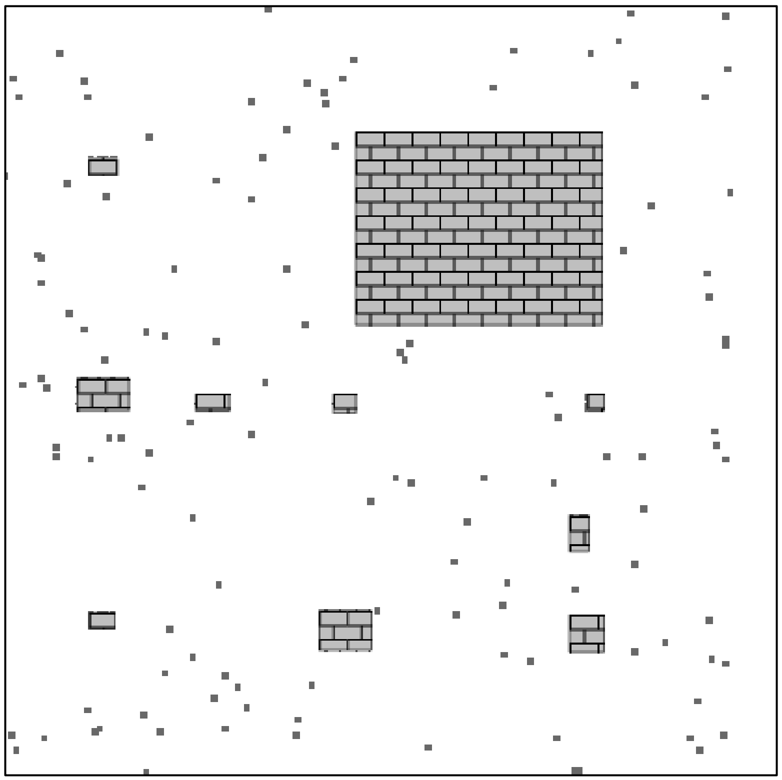

Uniform Sampling

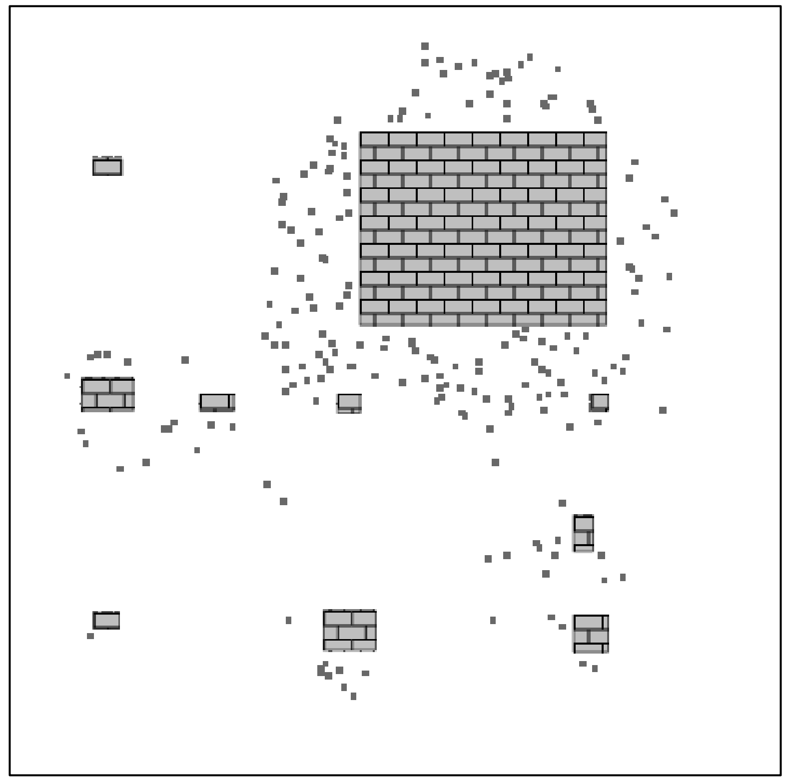

Sensor-Uncertainty Sampling

Running Time

	$tr(\Sigma)$	Build time	Search time
PRM	16.046	0.036	.001
BRM, Uniform Sampling	4.223	18.920	0.039
BRM, Sensor-Uncertainty Sampling	1.094	25.589	0.032

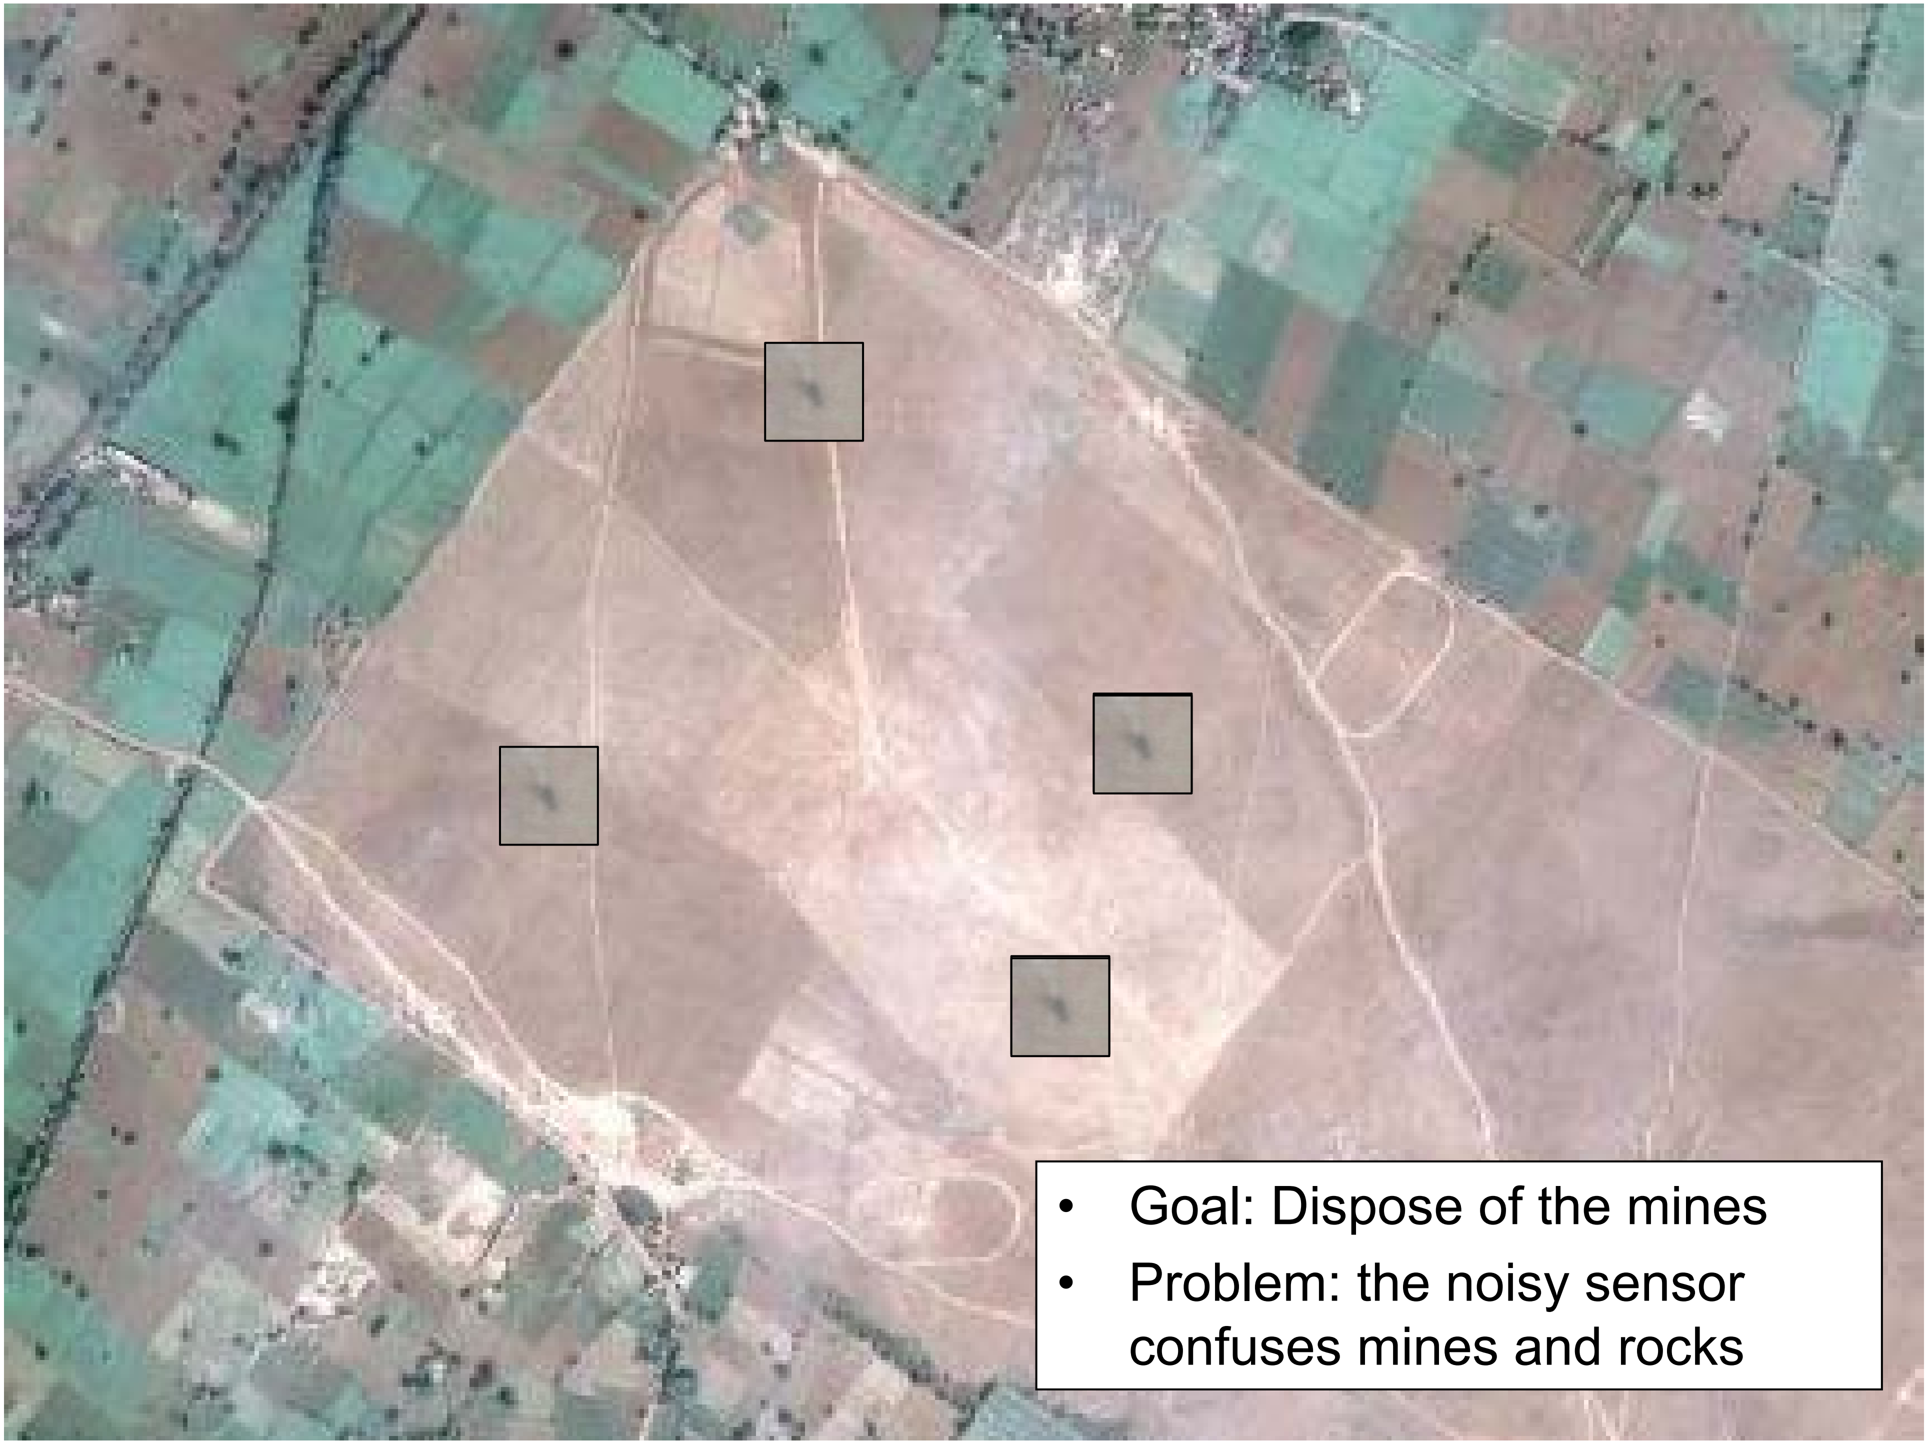

- Goal: Dispose of the mines
- Problem: the noisy sensor confuses mines and rocks

RockSample

- Given cost of flight, reward of disposing of actual mines…
- Search for a sequence of paths through the graph that maximize expected reward

- Problem: distribution is multinomial 
- Problem: posterior distribution is not deterministic

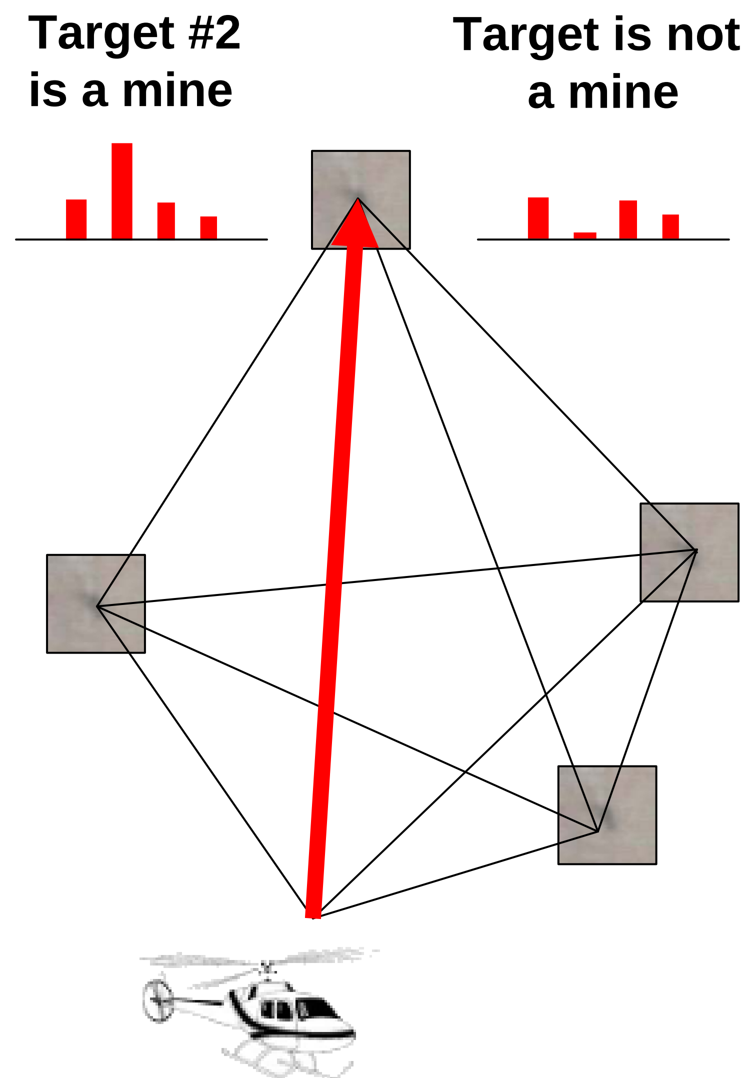

Posterior Belief Distribution

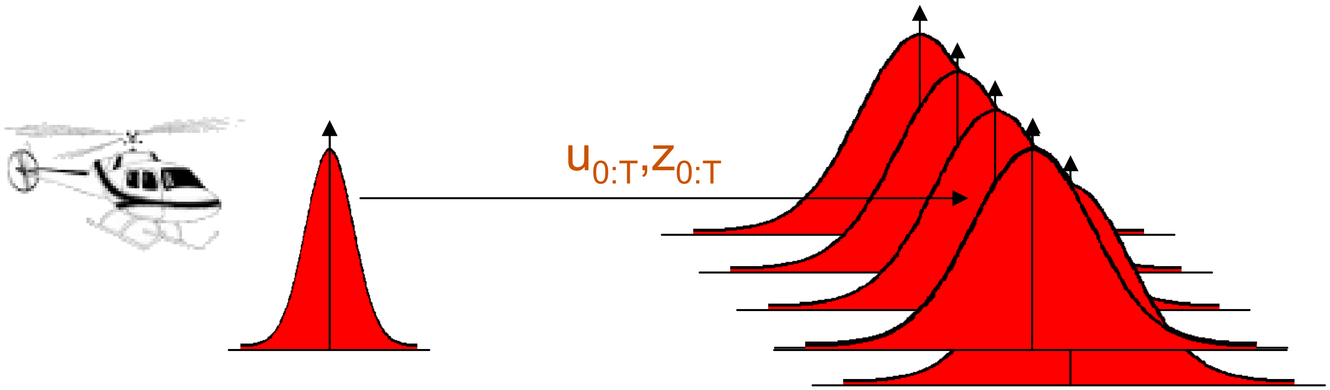

- Posterior belief no longer deterministic
- Action sequence leads to a distribution over posterior beliefs
- Compute expected reward over distribution of distributions
- Compare ∫R(b'|u'0:T)db' > ∫R(b|u0:T)db

- Analytic solution exists for linear Gaussian systems : O(n)
  - Approximate version available for exponential family distributions i.e., Poisson, Bernoulli, multinomial, Dirichlet, etc.

RockSample

- Search for a sequence of paths through the graph that maximize expected reward…
- Problem: graph may not contain optimal trajectory
  - Iteratively refine graph
  - Provable convergence to bounded optimal policy
- Planning under Uncertainty with Macro-Actions (PUMA)

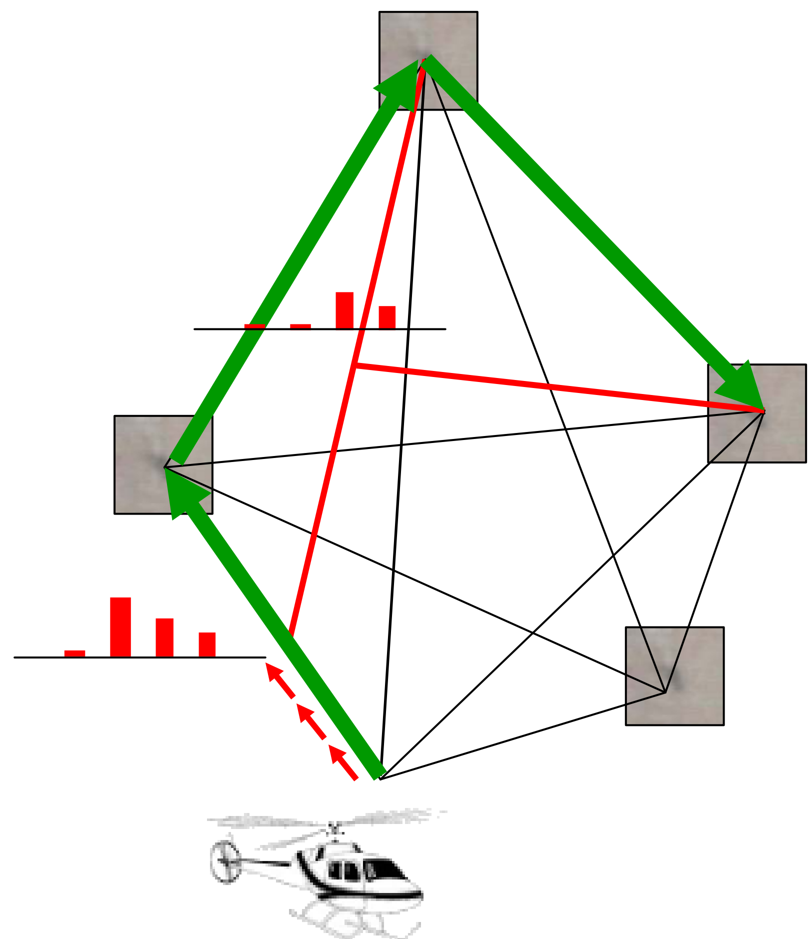

Experimental Performance

Problem	Algorithm	Ave. rewards		Online time(s)	Offline time (s)
ISRS (8,5)	SARSOP	12.10	±0.26	0.00	10000
	Naïve FS	9.56	±1.08	3.36	0.00
	Hand-coded SCP	19.71	±0.63	0.74	0.00
	PUMA	17.38	± 1.41	162.48	0.00

- ISRS: 2048 states
- Largest version of this problem solved so far: $10^4 \times 2^{30}$ states, 1-2 minutes per step

Model-uncertainty Planning

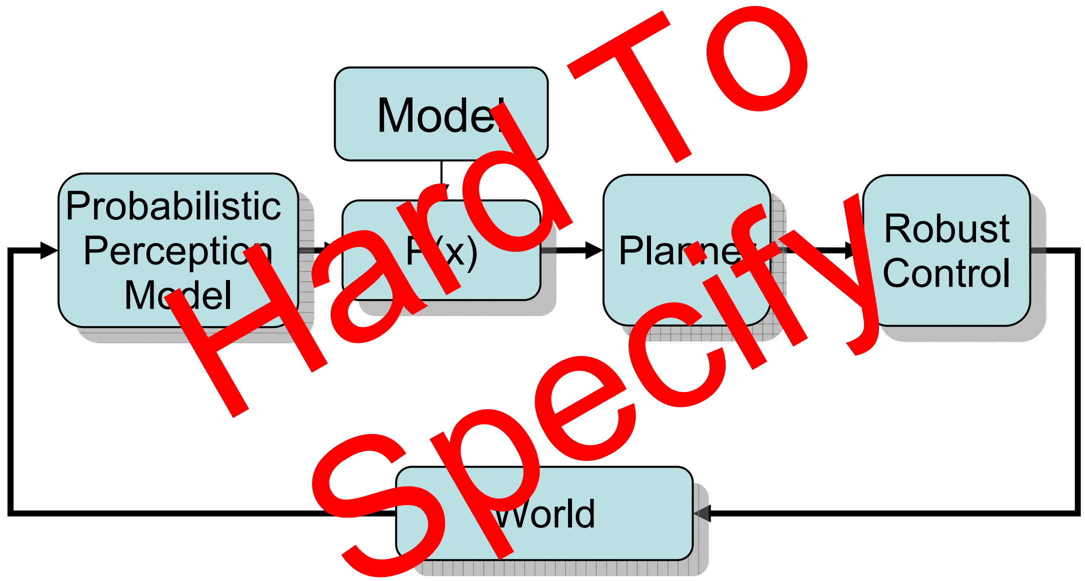

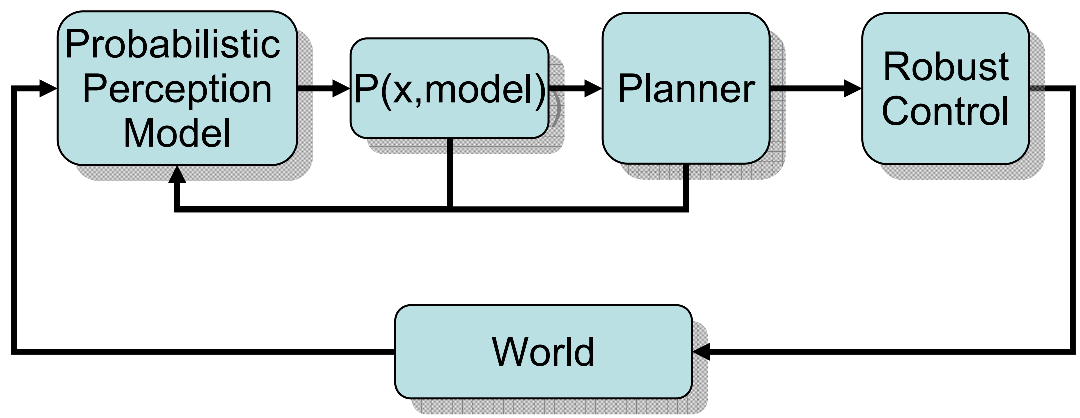

Information Dynamics

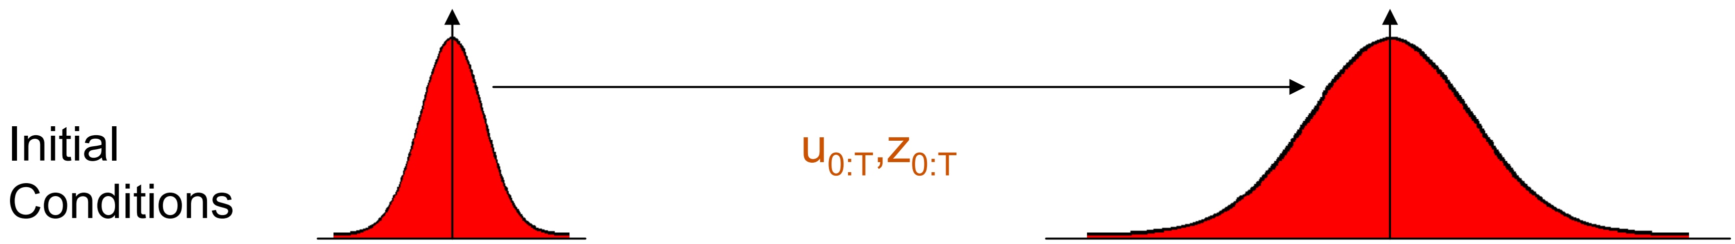

- How do I predict the posterior efficiently?
  - Choose the right matrix inversion lemma
- Can I make the posterior small efficiently without explicit prediction?
  - Machine learning

Reinforcement Learning

- Input:
  - Current and goal vehicle pose
  - Current map estimate
- Output
  - Trajectory that minimizes expected cost

- Learn actions that minimize expected cost in practice

- Core algorithm: stochastic function approximation

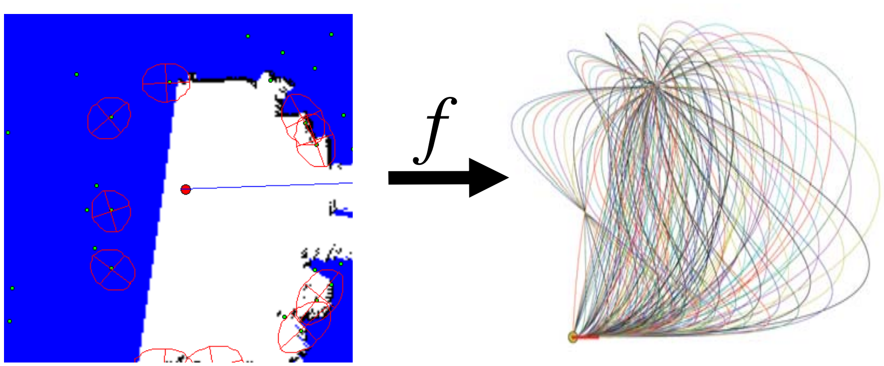

$$f : s \times M \times g \to a$$

$$R^3 \quad \mathbf{R^n} \quad R^3$$

Map Error Minimization

Shortest Path Explorer

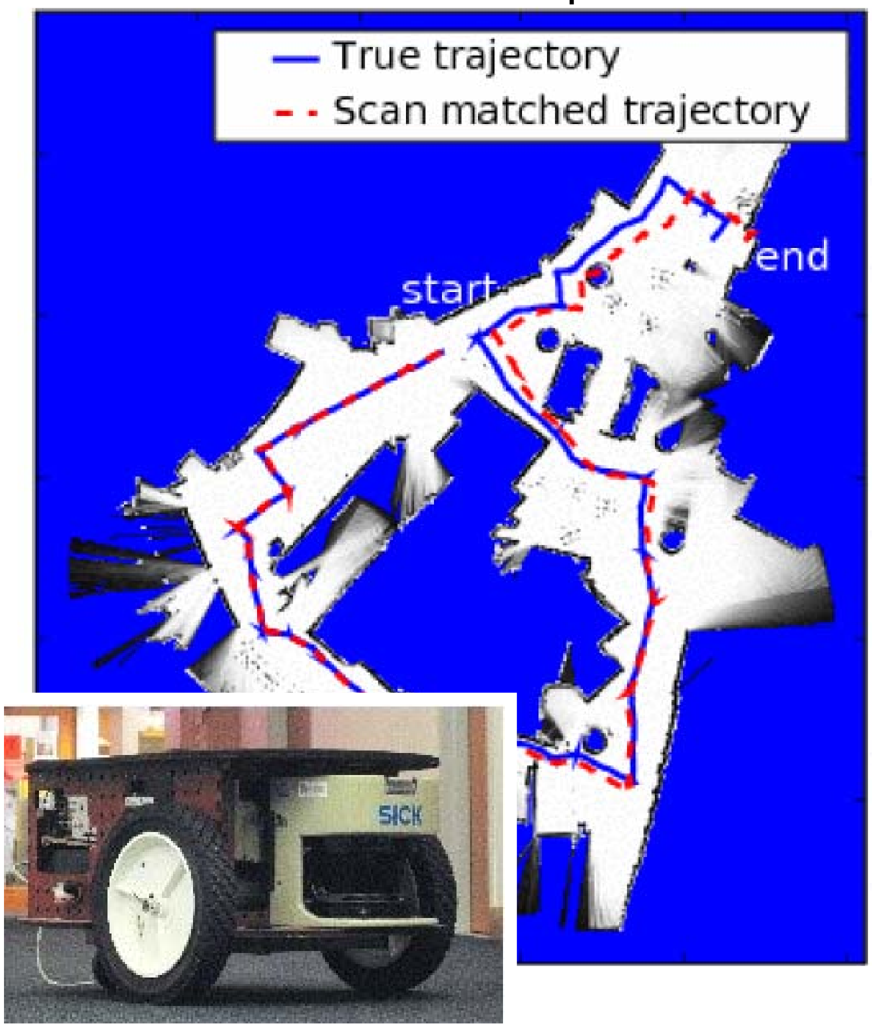

Learned Explorer

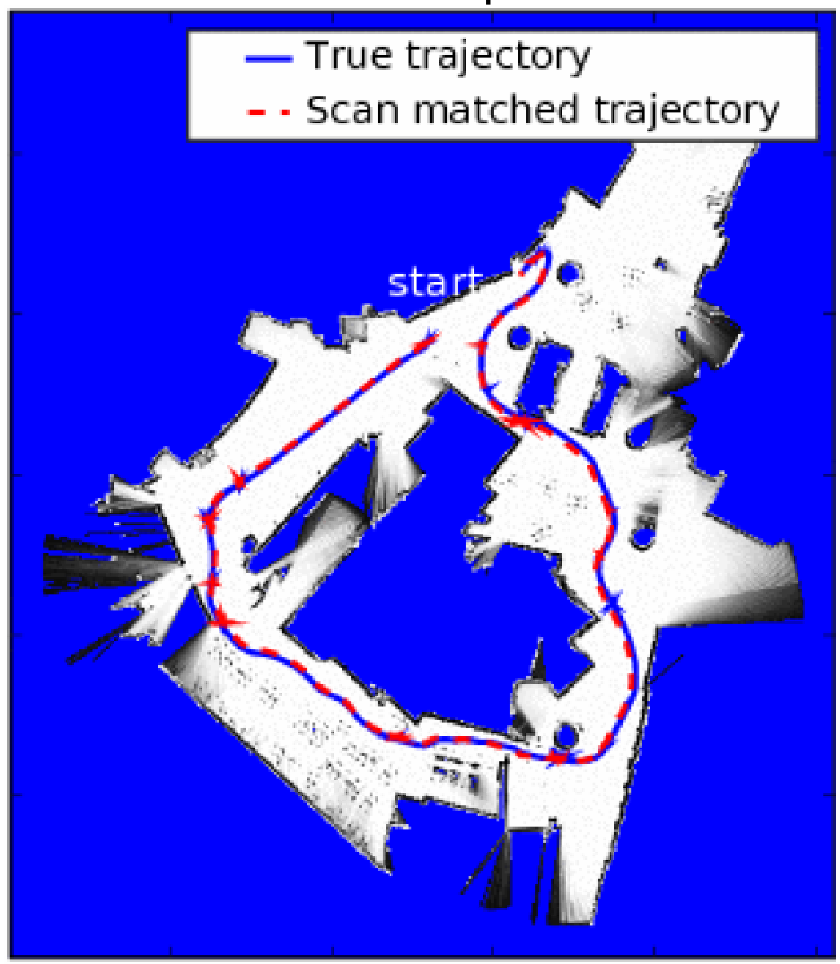

Map Error Minimization

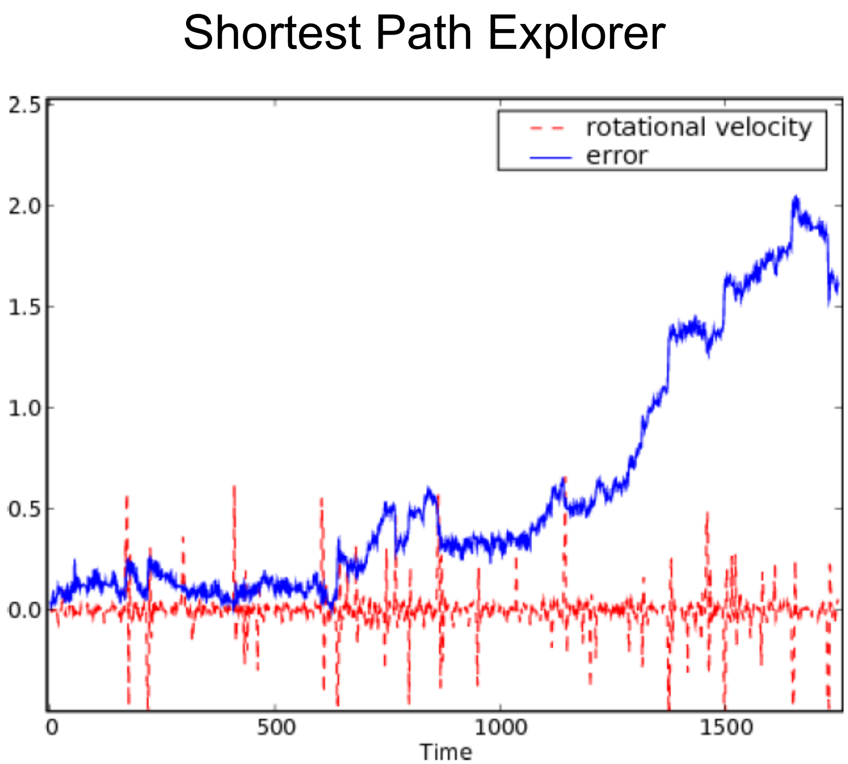

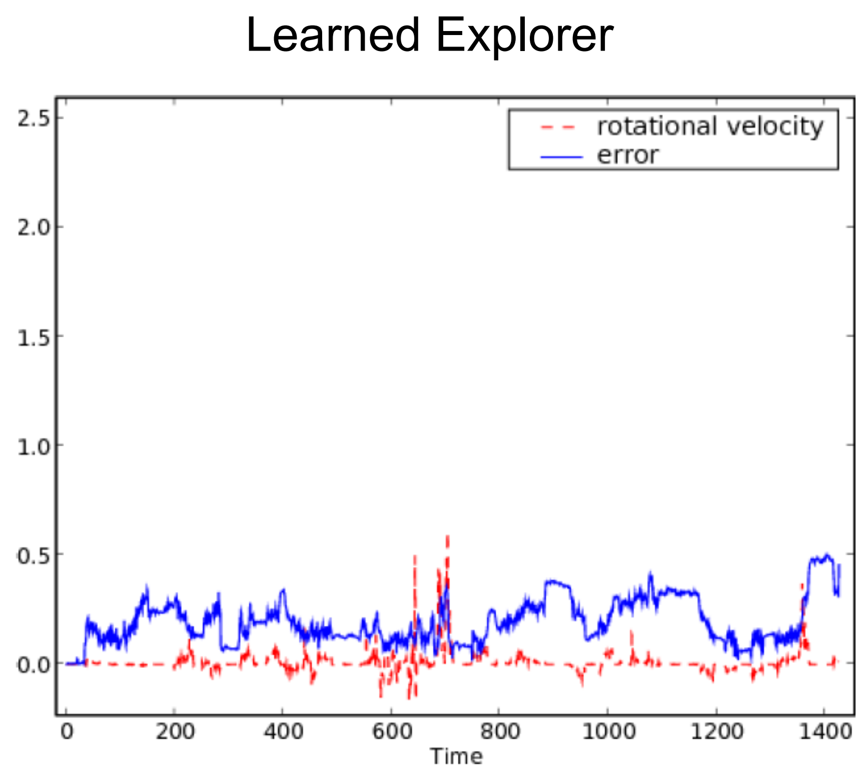

Previous state of the art: computing each exploration action takes 30 minutes

Learned controller: computing each exploration action takes milliseconds


Autonomous
Navigation

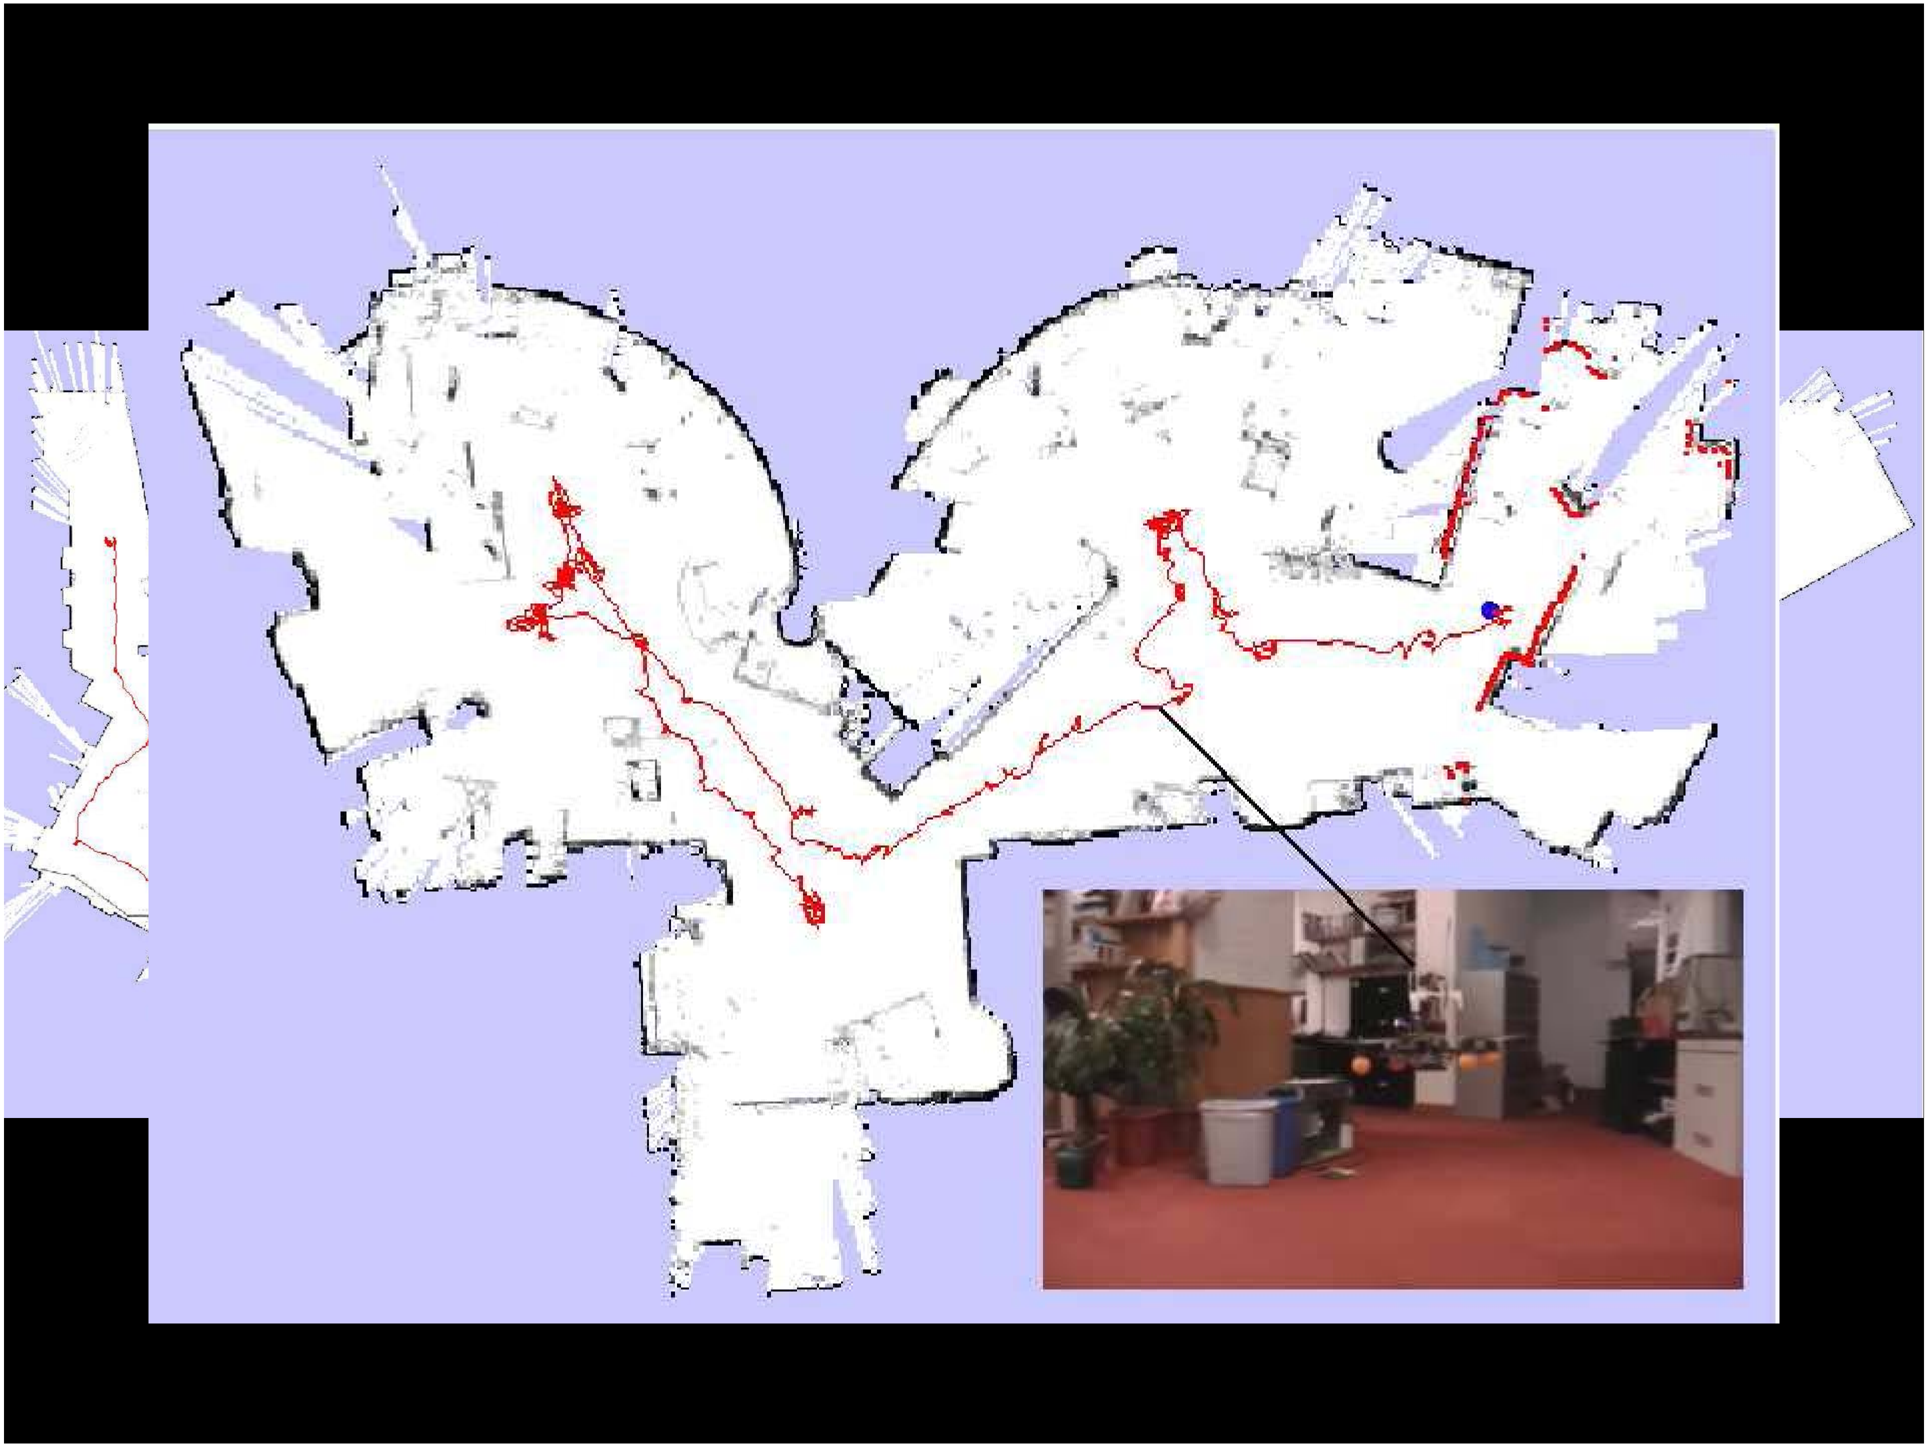

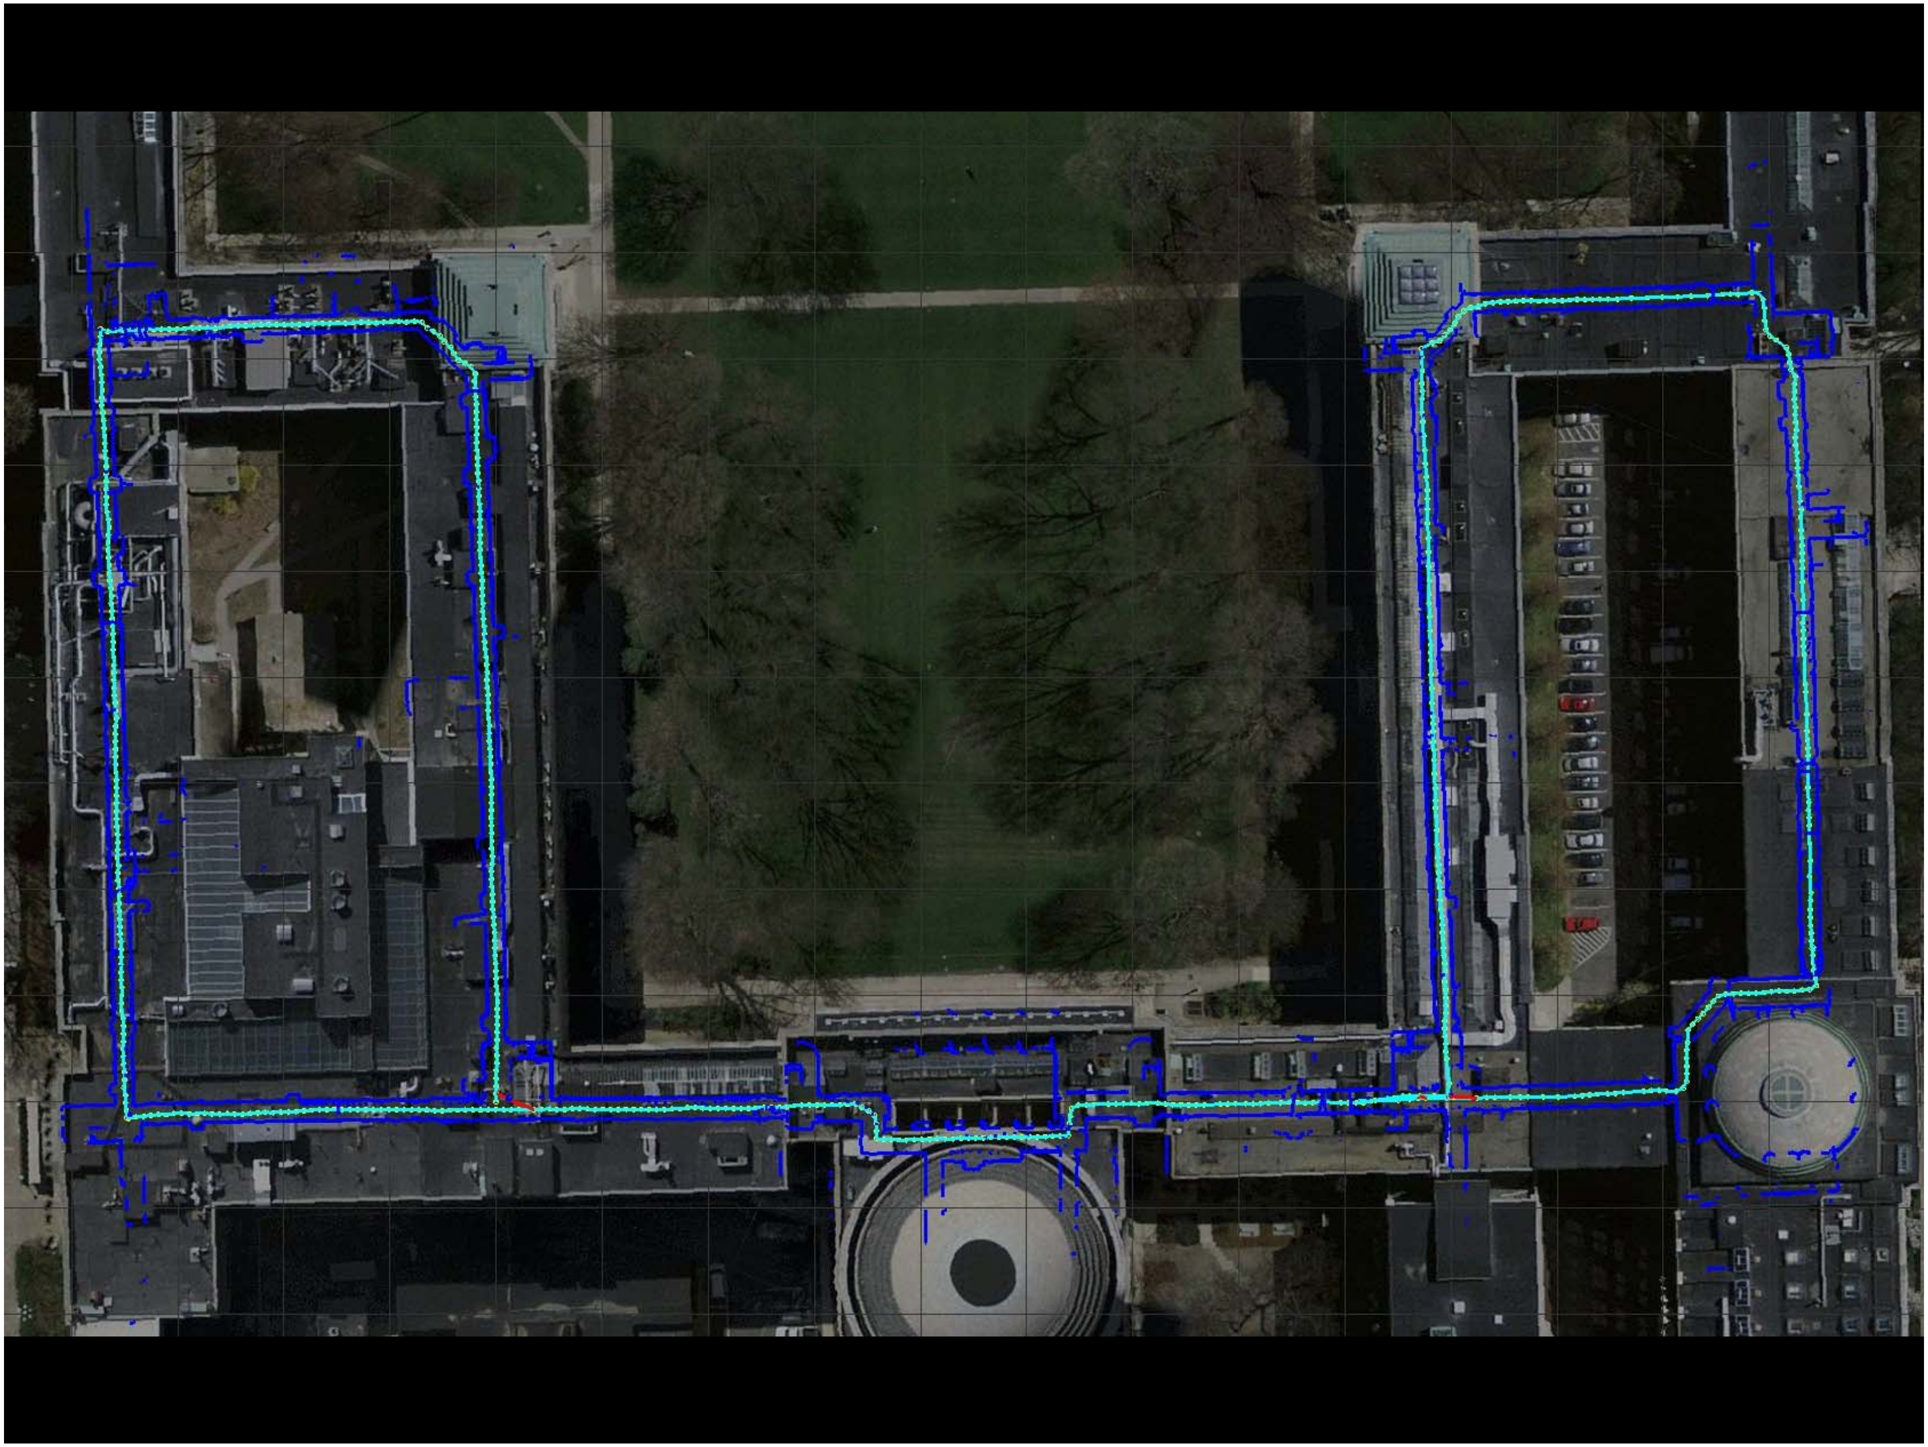

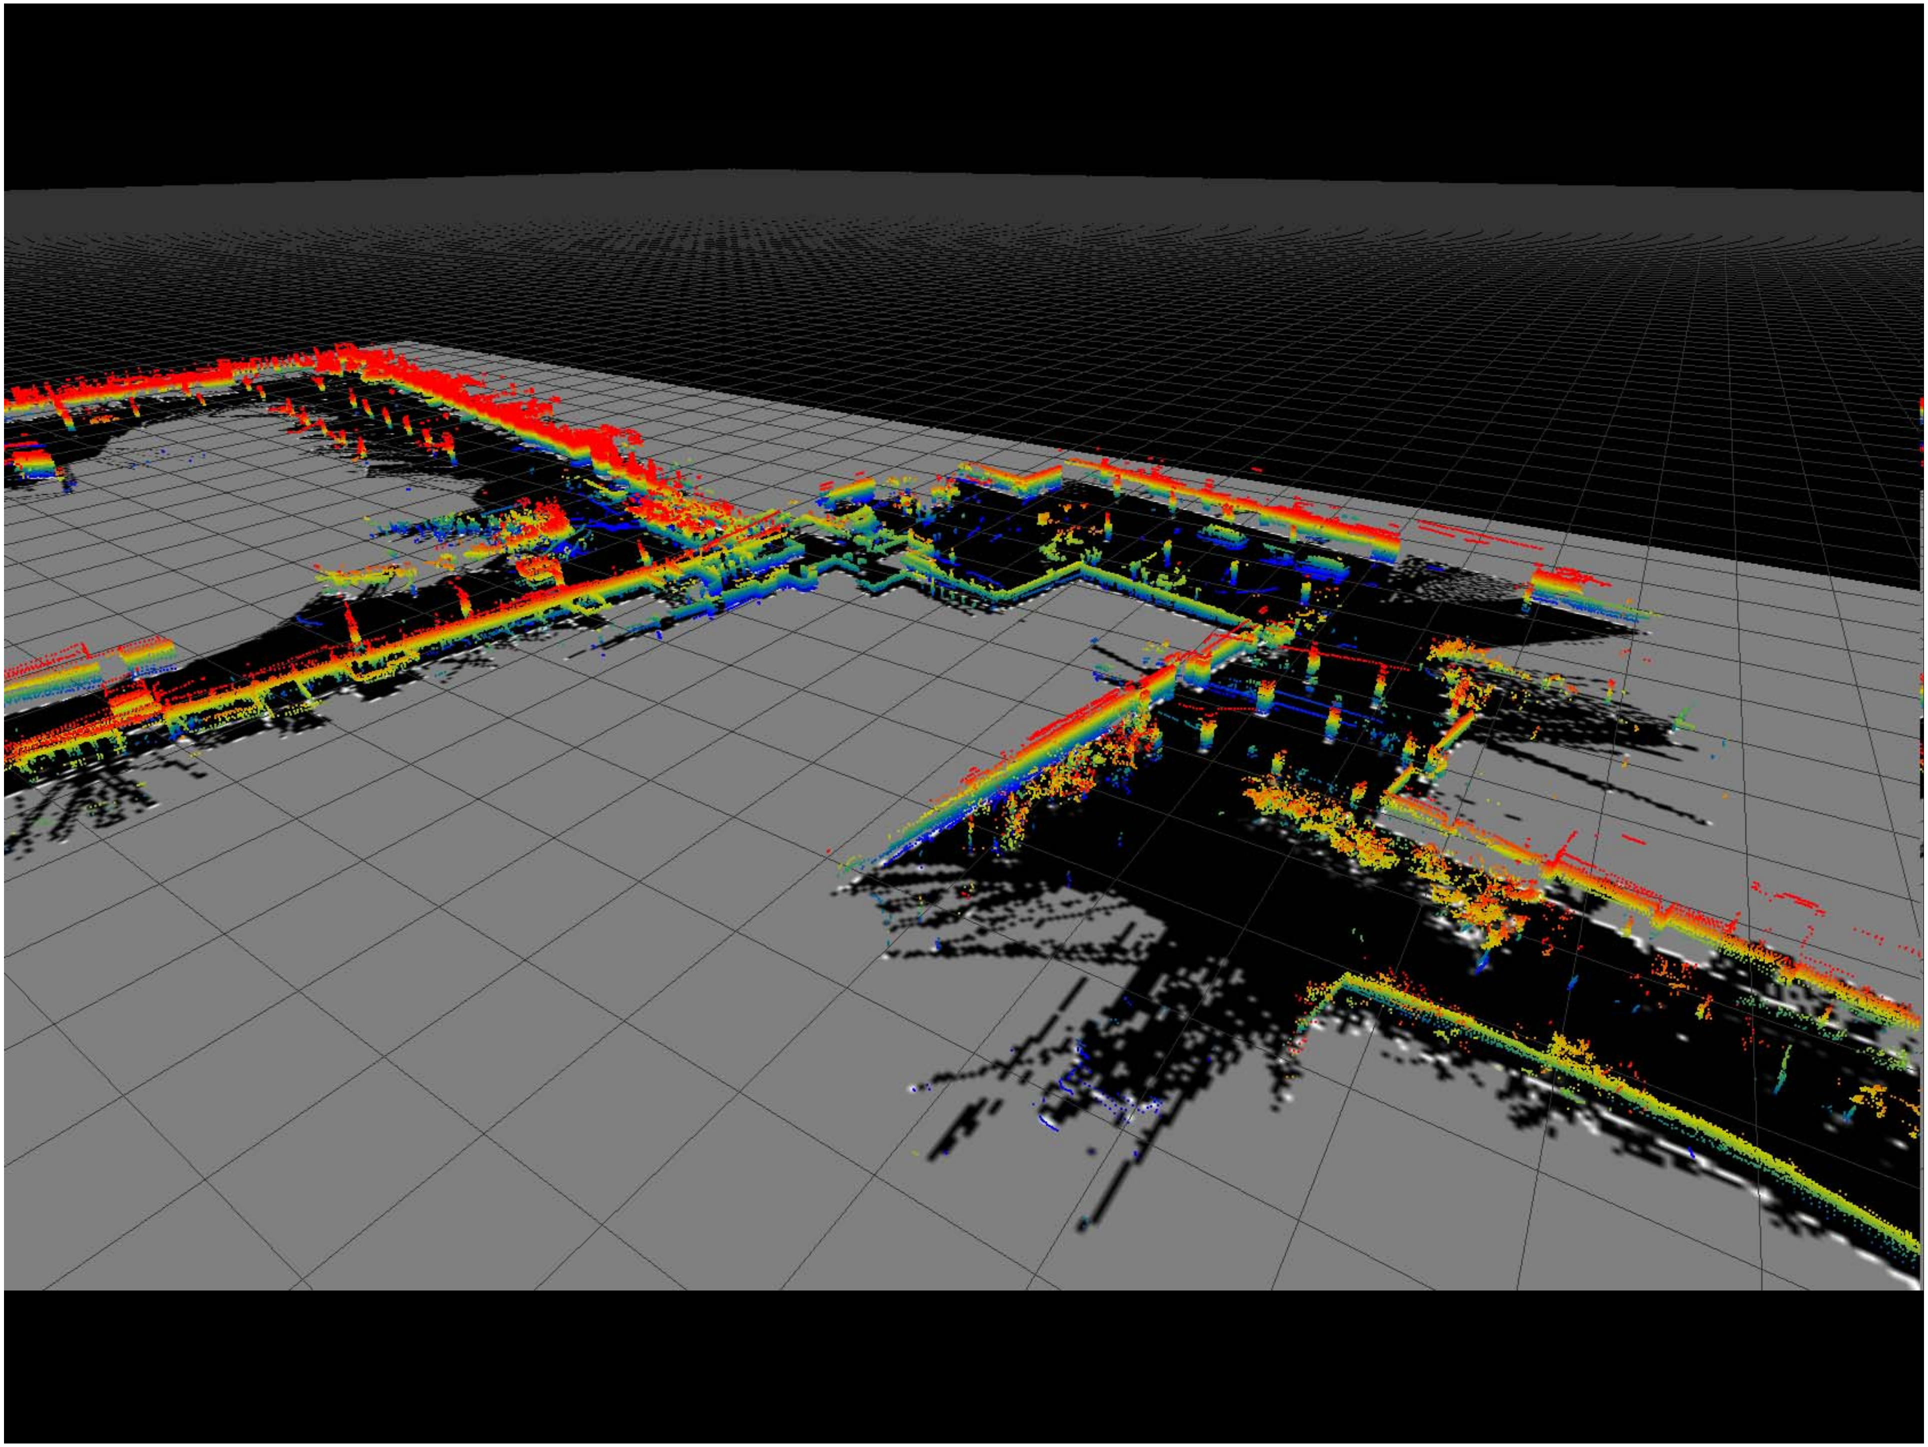

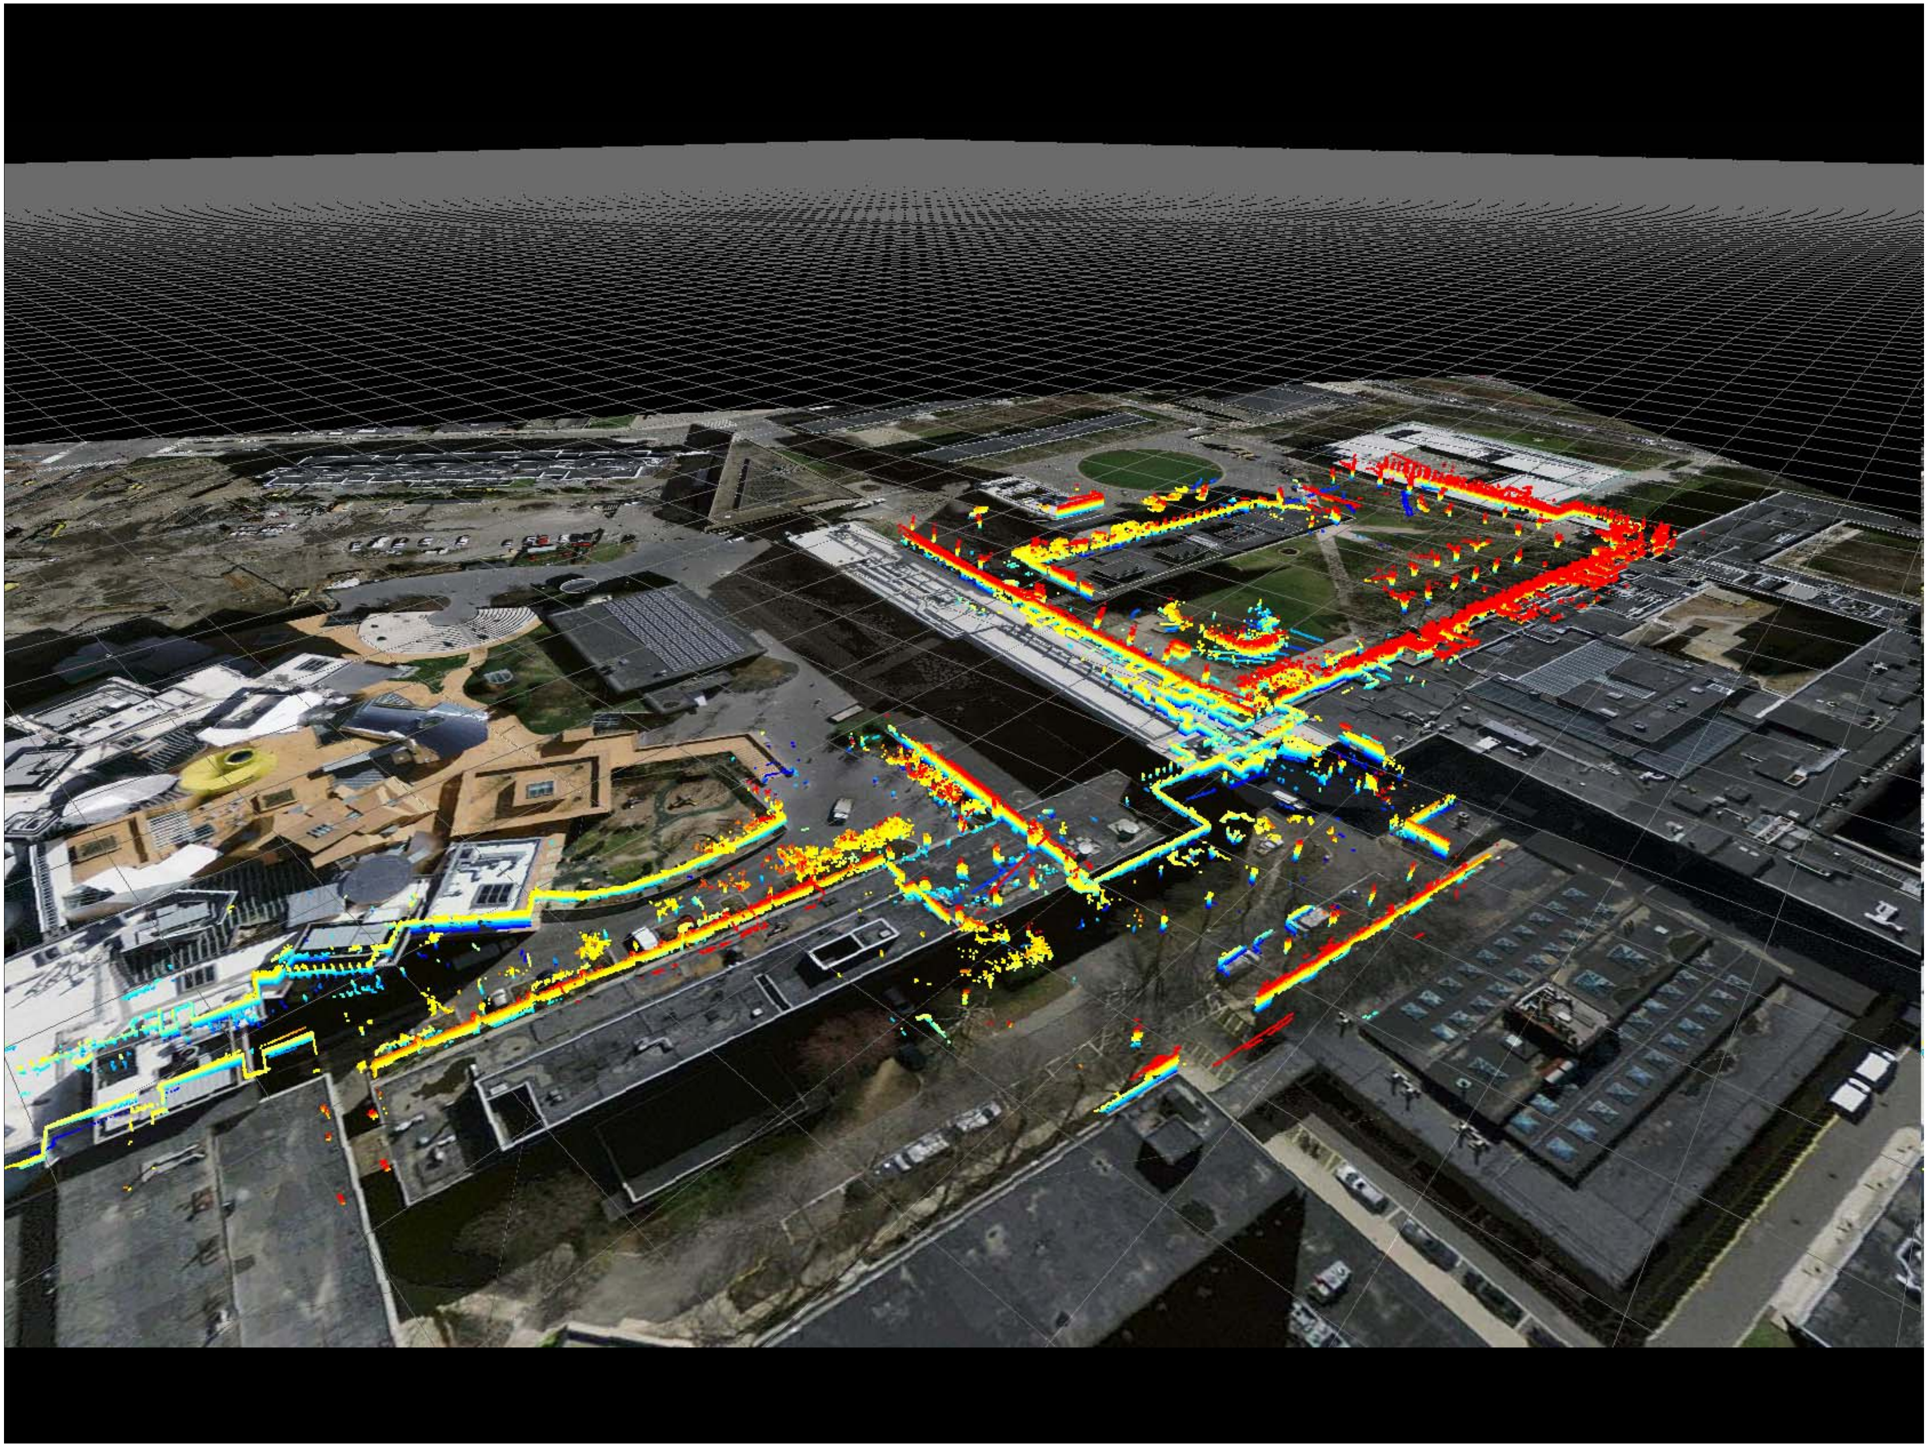


IARC
IARC
5TH MISSION
IARC
5

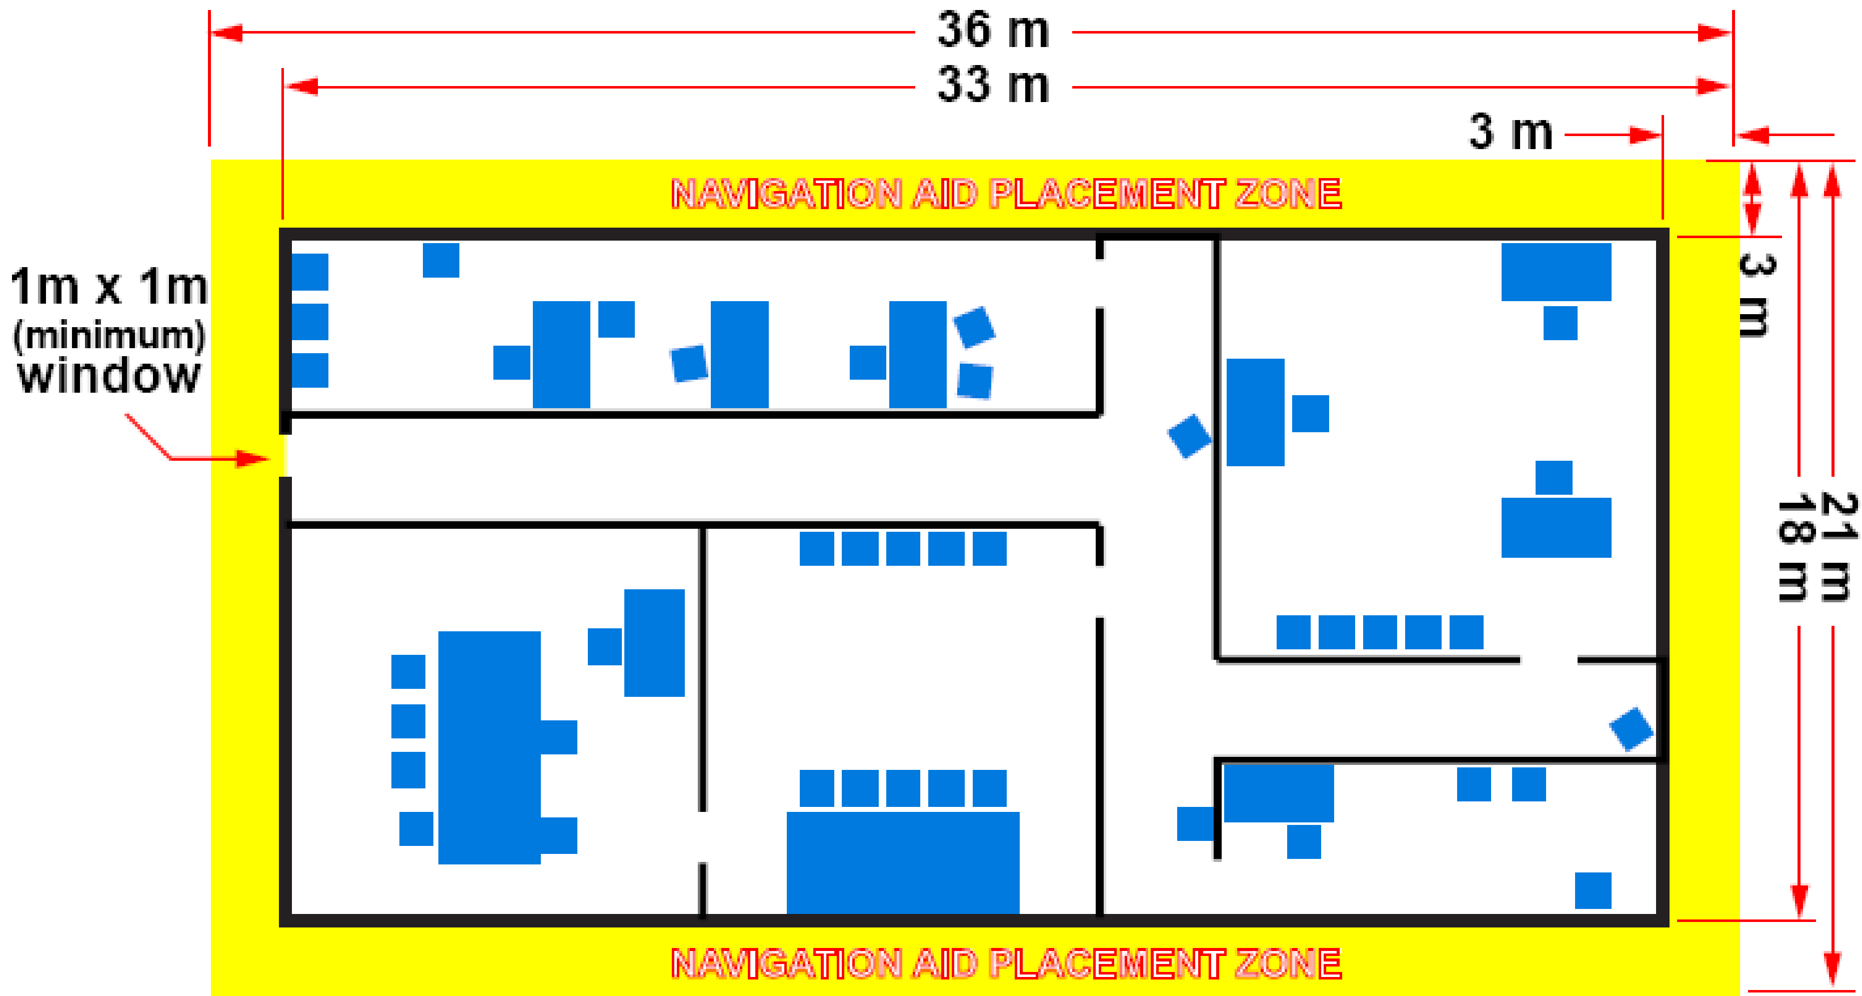

NOTE: Internal wall and obstacle placement is purely notional.
Actual placements and numbers of rooms may differ.
No entry door will be less than 1 meter in width.


COLEGIO
C
A
A

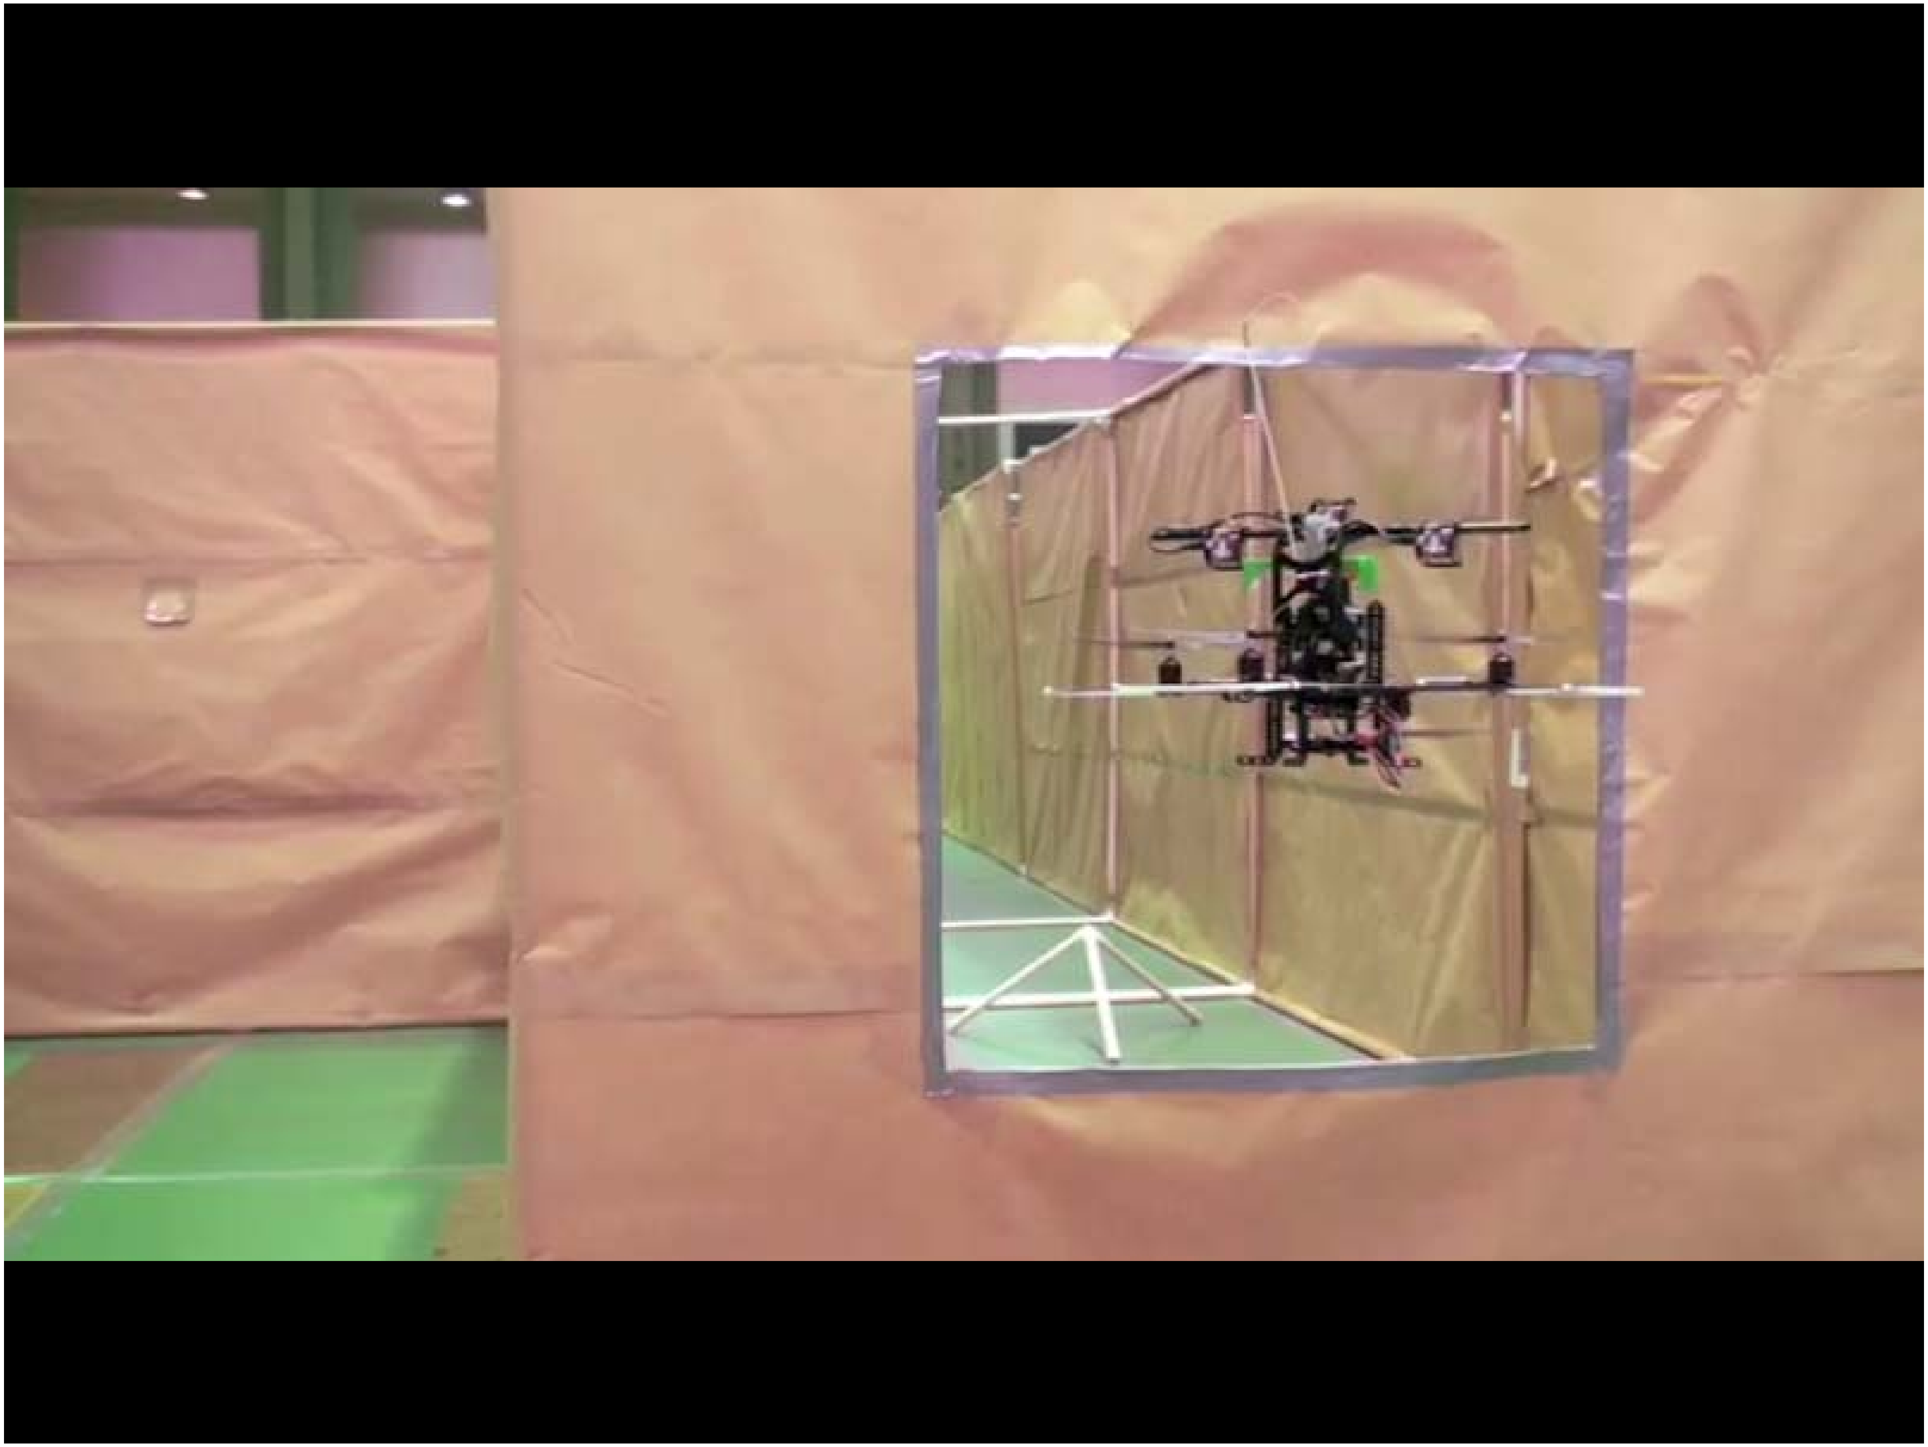

Autonomous Entry

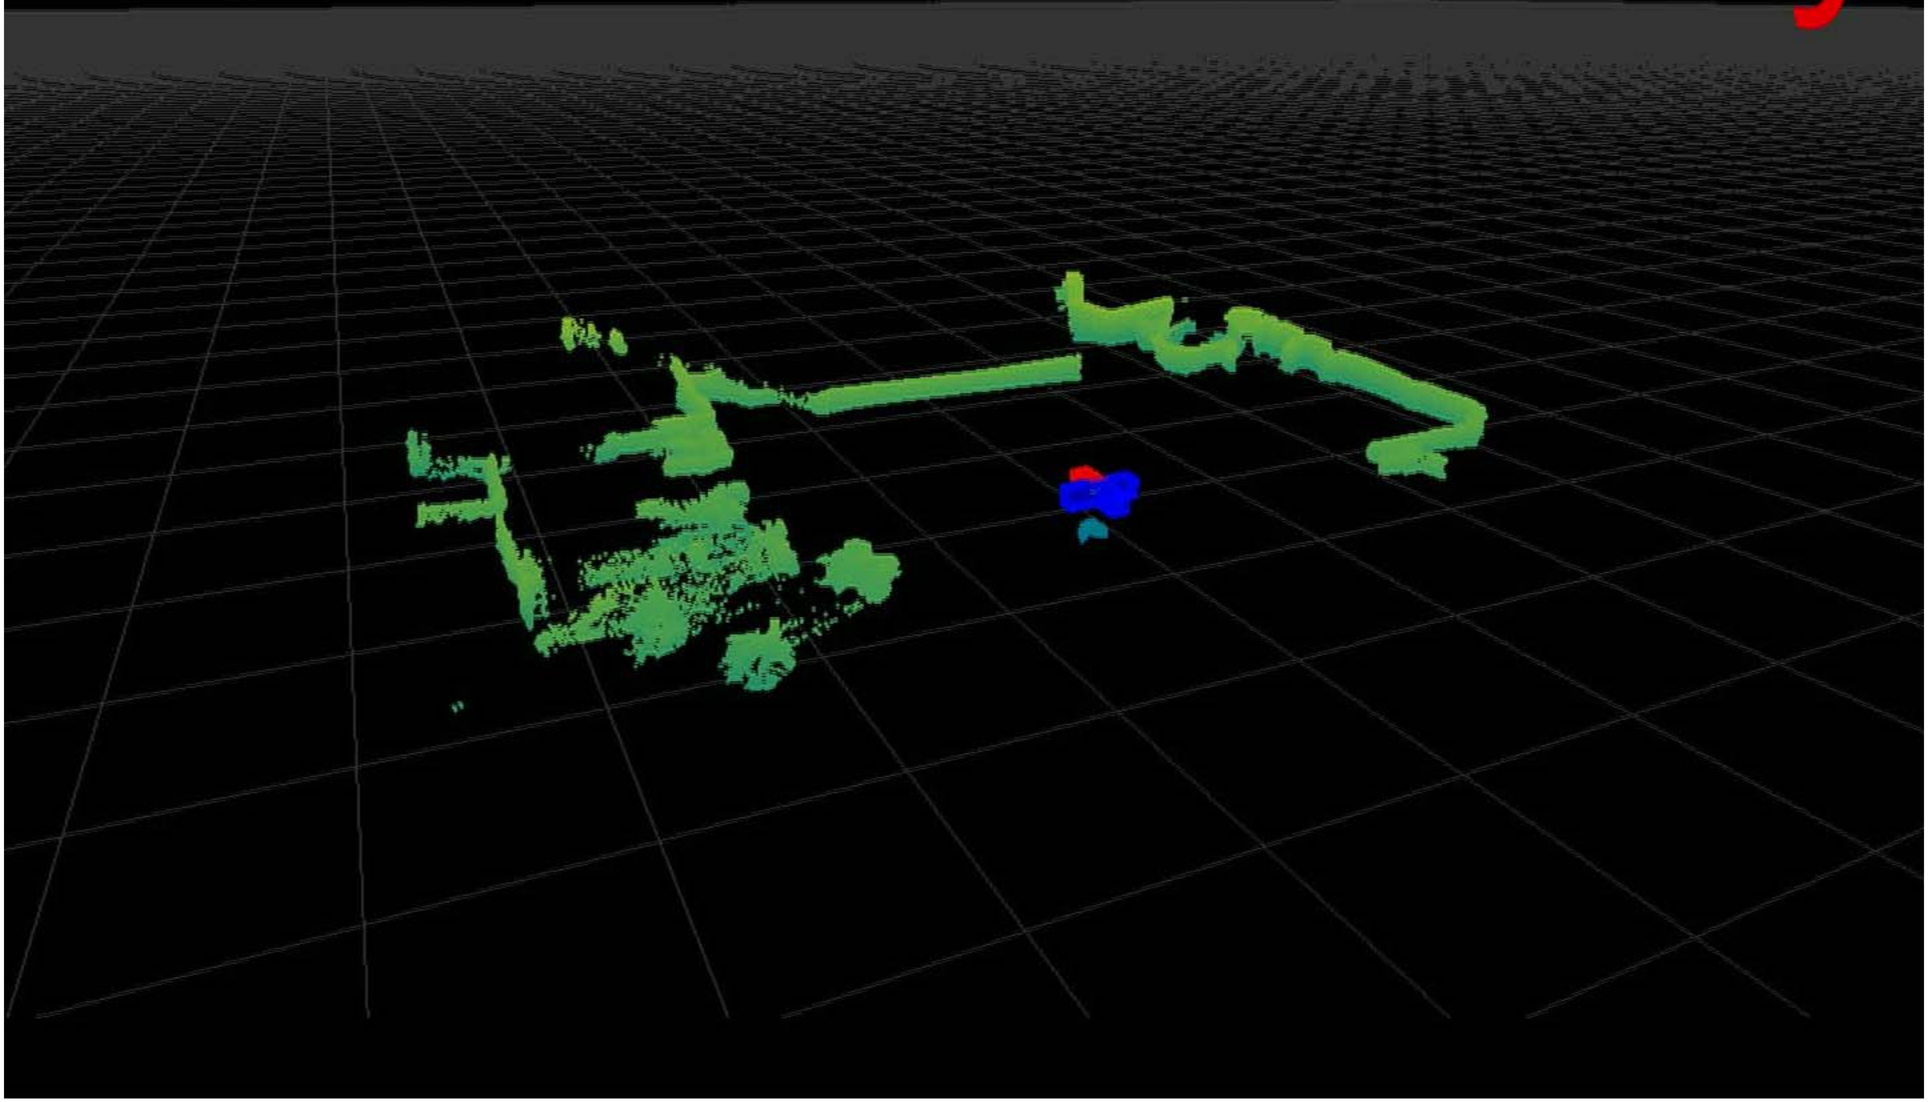

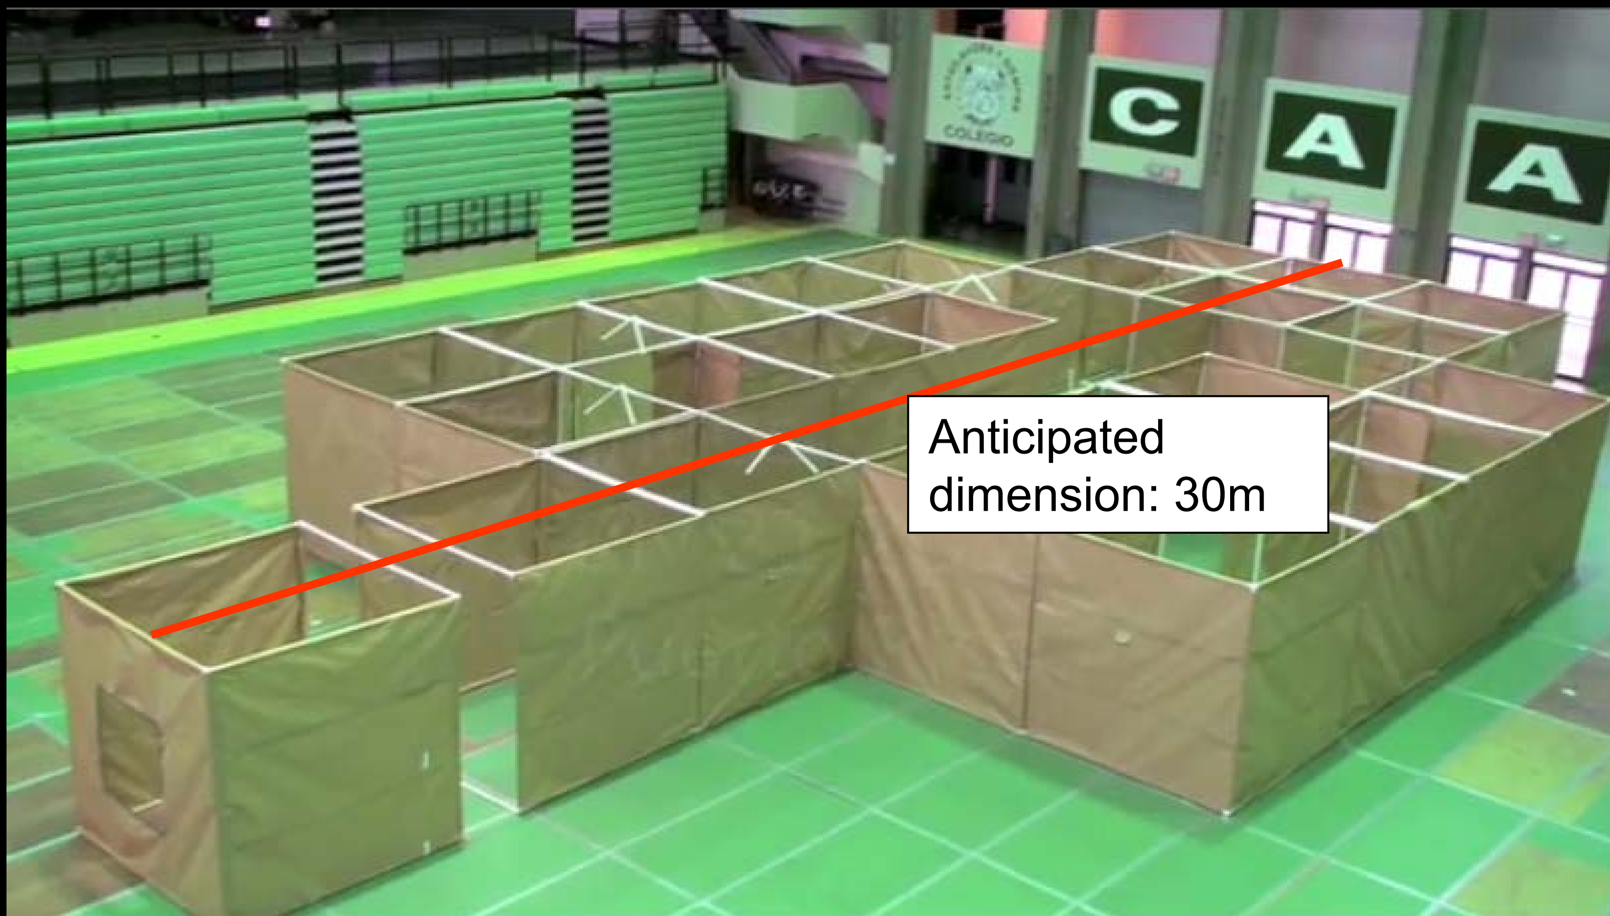
Anticipated
dimension: 30m
C
A
A
COLEGIO


UPR
UNIVERSIDAD DE PUERTO RICO
1903
RUM

[redacted]


Interactive

Summary

- Robust, long-term autonomy in large-scale environments
- Planning algorithms for worlds in which we have limited knowledge of the state, model of the system, or a map of the world
- Key Issue: Control of Information
- Technical approaches:
  - Understanding how information propagates
  - Machine learning