Hierarchical Multi-Objective Planning For Autonomous Vehicles

Alberto Speranzon
United Technologies Research Center

UTC Institute for Advanced Systems Engineering Seminar Series

Acknowledgements and References

- Shaunak Bopardikar – UTRC Berkeley
- Dennis Ding – UTRC East Hartford
- Brendan Englot – UTRC East Hartford
- Alessandro Pinto – UTRC Berkeley
- Amit Surana – UTRC East Hartford

- X. Ding, B. Englot, A. Pinto, A. Speranzon and A. Surana, "Hierarchical Multi-objective Planning: From Mission Specifications to Contingency Management", To be published at ICRA 2014

- S. Bopardikar, B. Englot and A. Speranzon, "Multi-Objective Path Planning in GPS Denied Environments under Localization Constraints", To be published at ACC 2014

- S. Bopardikar, B. Englot and A. Speranzon, "Robust Belief Roadmap: Planning Under Uncertain And Intermittent Sensing", To be published at ICRA 2014

- X. Ding, A. Pinto and A. Surana, "Strategic Planning under Uncertainties via Constrained Markov Decision Processes", Appeared in ICRA 2013

Problem: High-Level Mission Specifications

Autonomous missions in uncertain environments require:

1) Support optimization over multiple costs
2) Handle logical/spatial/temporal constraints
3) Deal with contingencies at multiple temporal and spatial scales

Mission (example):

Starting from **START**, go to **PICKUP** location, then go one of the **DROPOFF** locations before heading back to **START**. Minimize the expected time of arrival with the constraints that the mission can be accomplished with at least 60% probability and total threat exposure is less than 0.4

Mission + Motion Planning

Mission VS Motion Planning

"Starting from **START**, go to **PICKUP** location, then go one of the **DROPOFF** locations before heading back to **START**. Minimize the expected time of arrival with the constraints that the mission can be accomplished with at least 60% probability and total threat exposure is less than 0.4"

Mission level planning:

- Reach some locations (**START PICKUP DROPOFF**)
- Optimize a primary goal (expected time) and satisfy constraints (probability of mission success and threat exposure)

Motion level planning:

- "Figure out" how to do execute the above in a complex city-like environment flying low between buildings to keep coverage
- Ensure that you are generating trajectories that are compatible with the underlying vehicle dynamics

World Model for Mission

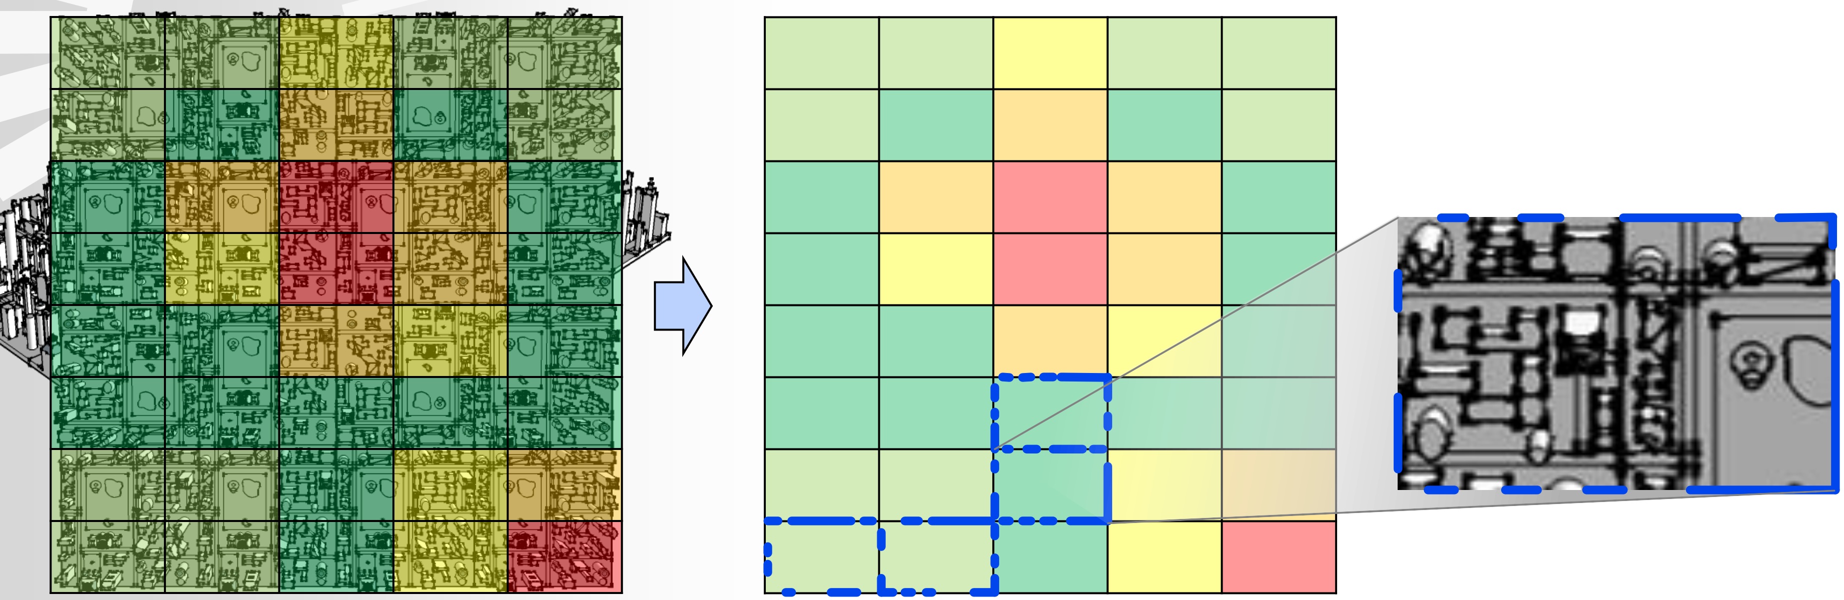

Labelled Markov Decision Process

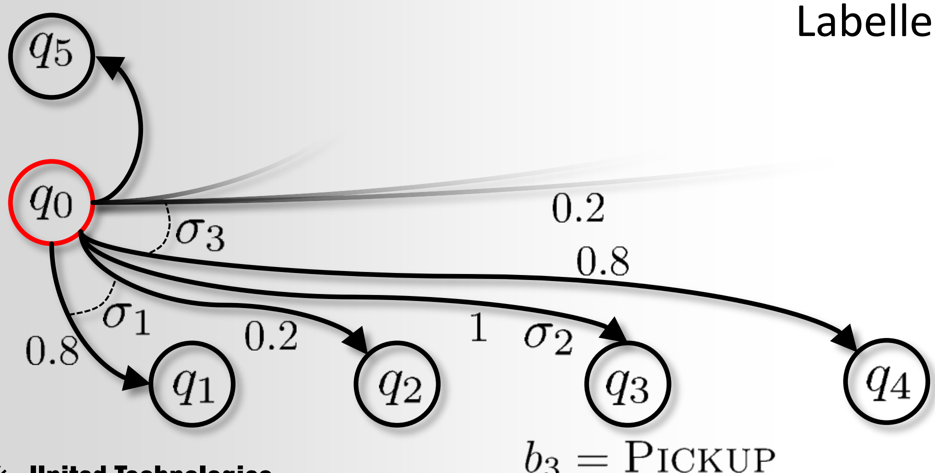

q_i is the state (facet, orientation)

σ_i is the action

$P(q_i, \sigma_i, q_j)$ is the probability of transition

b_i label at state i

$g_i(s_i, \sigma_i)$ is the cost of the action σ_i

Mission Level Planning

- Given a mission specification expressed as linear temporal logic (LTL) obtain Deterministic Finite State Automaton (DFA)

Starting from **START**, go to **PICKUP** location, then go one of the **DROPOFF** locations before heading back to **START**

$$\phi = \neg \texttt{Fail}\, \mathsf{U} \Big(\mathsf{F}\, \texttt{Start} \wedge \mathsf{F}\, \texttt{Pickup} \wedge \mathsf{F}\, \texttt{Dropoff}$$
$$\wedge\, \mathsf{G}\, (\texttt{Start} \rightarrow \mathsf{X}\,\mathsf{F}\, \texttt{Pickup} \wedge \texttt{Dropoff} \rightarrow \mathsf{X}\,\mathsf{F}\, \texttt{Start}) \Big)$$

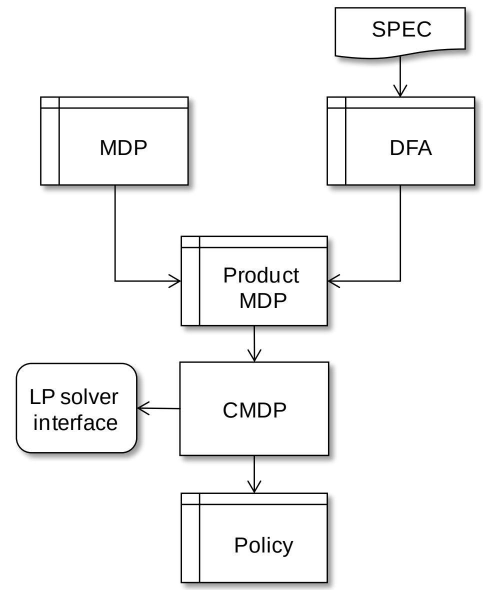

- MDP represents the world, the actions and the costs

- Combine the MDP and DFA to obtain a CMDP

- Solve CMDP to obtain (randomized) mission level policy (plan)

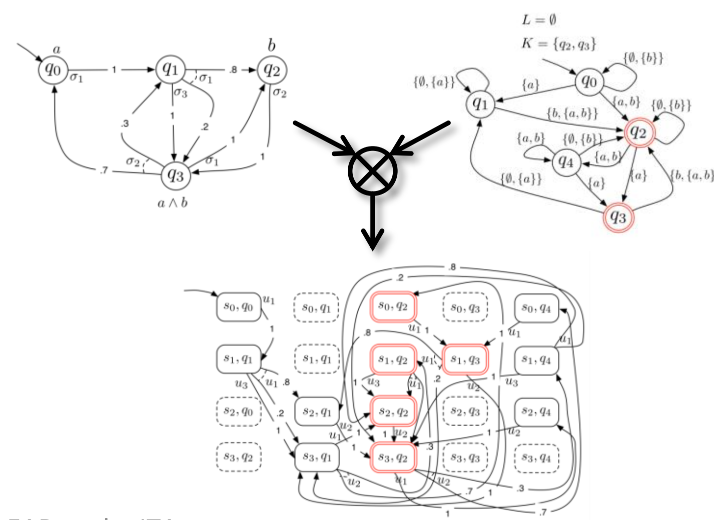

Motion Level Planning

- Responsible to execute the mission level policy at a lower level

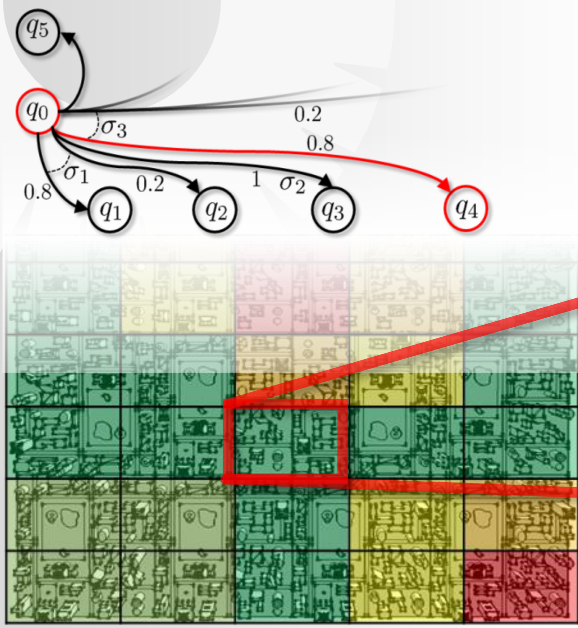

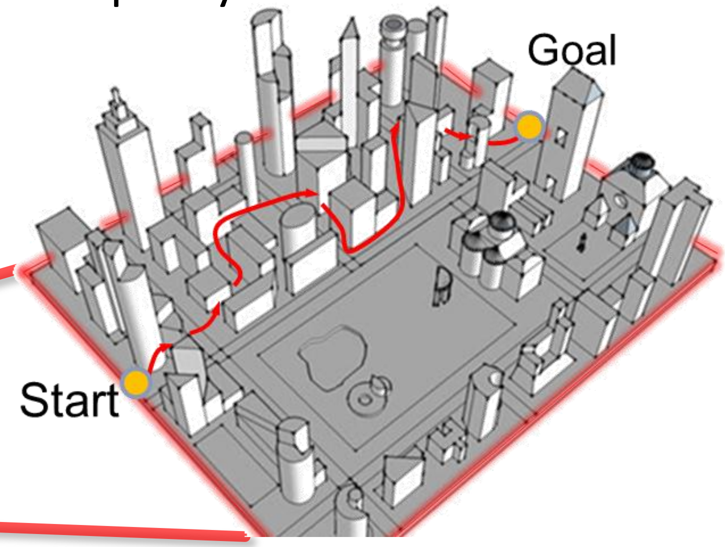

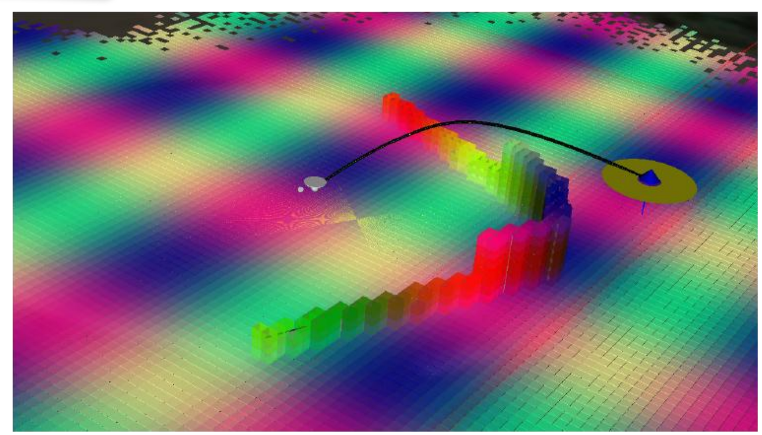

- Use of evidence grid to represent occupied/unoccupied space

- How do we ensure that there is *"consistency"* between the mission level planning cost and constraints and the low-level planning objective?

Hierarchical Planning

- Costs and constraints between the different levels of the hierarchy are in correspondence across layers

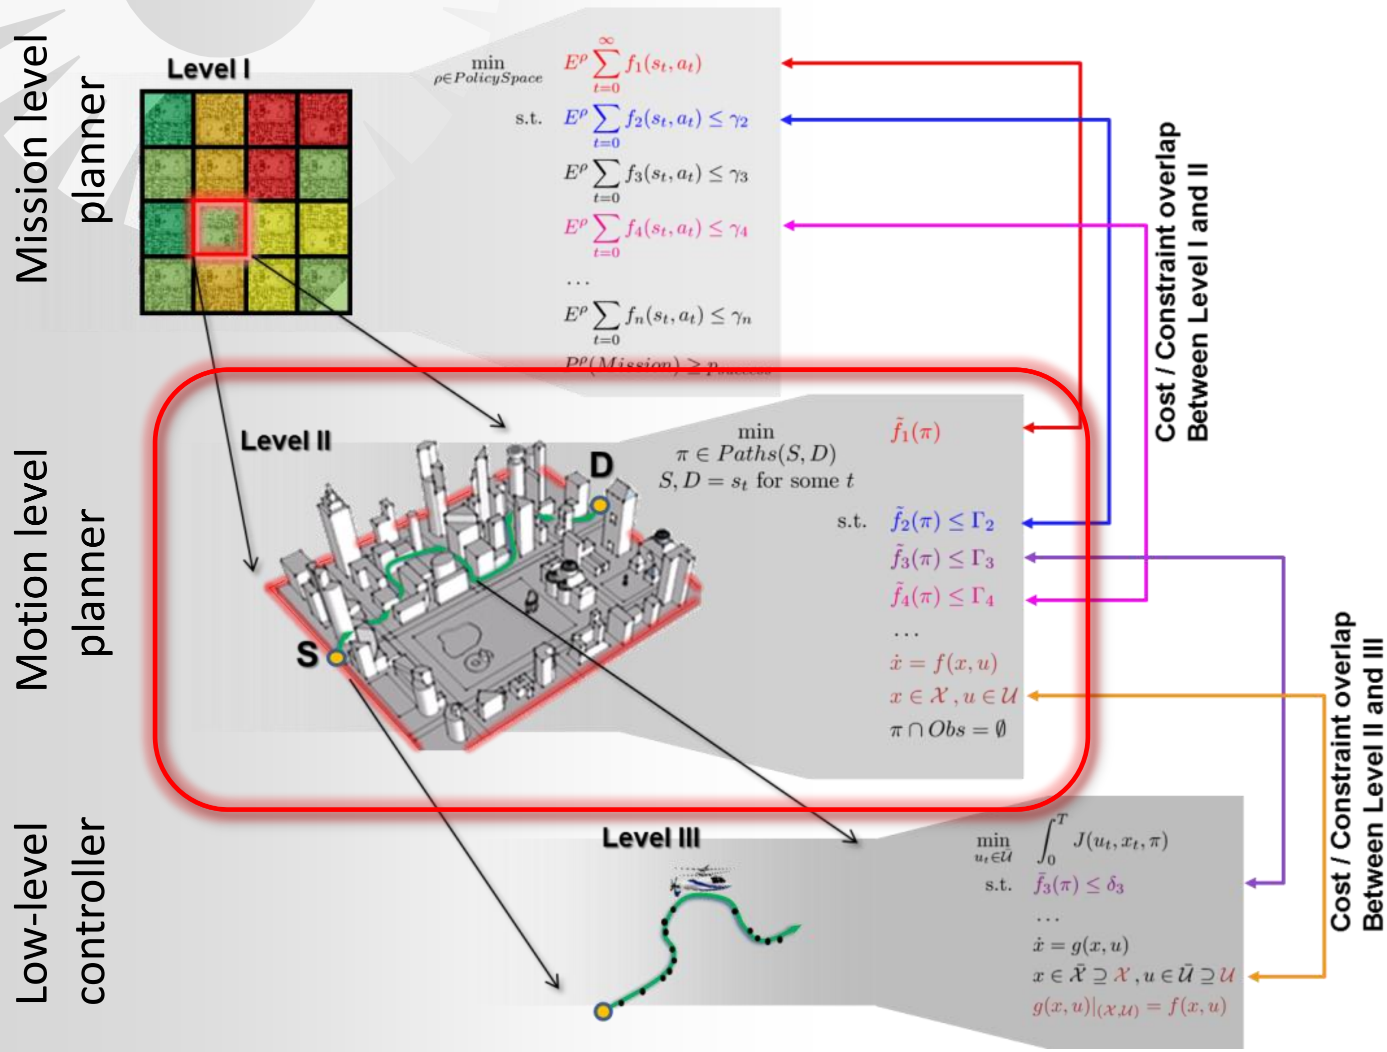

Probabilistic Roadmaps

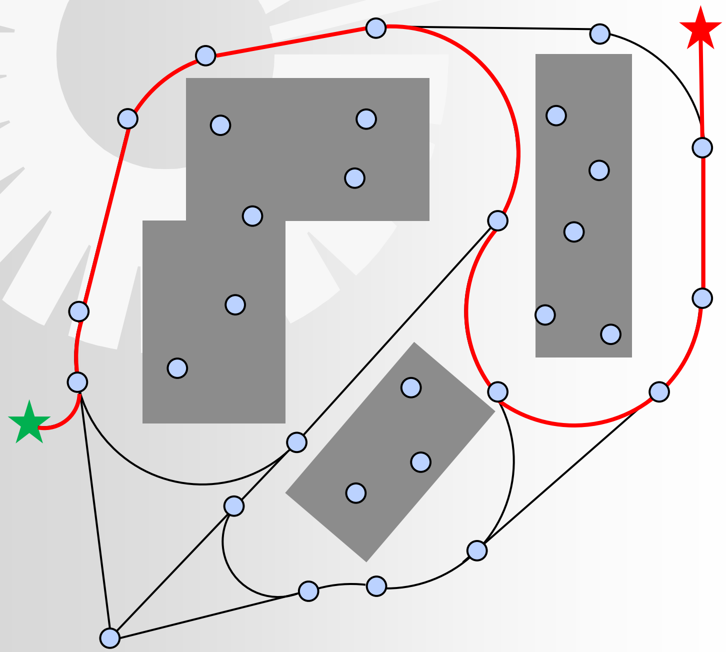

1. Randomly sample the configuration space
2. Remove samples that are not collision free
3. Determine path compatible with vehicle dynamics that connects the nodes
4. Connect Start and Goal to closest nodes

- Samples can be drawn in a deterministic or in a stochastic fashion
- Useful for planning in higher dimensional spaces - e.g. in 3D considering (x, y, θ) or 6D considering position $(x, y, z) +$ velocity (v_x, v_y, v_x)
- PRM sampling methods are probabilistic complete

Multi-objective Path Planning

- We are interested to compute a plan that minimizes two costs functions $C(.)$ and $Q(.)$
- To pose this problem we consider the cost function $C(.)$ as primary cost and $Q(.)$ as a secondary cost (constrains) and pose the following problem where now b is considered a free variable

$$\min_{\pi \in \mathcal{R}} \quad C(\pi)$$
$$\text{s.t.} \quad Q(\pi) \leq b$$

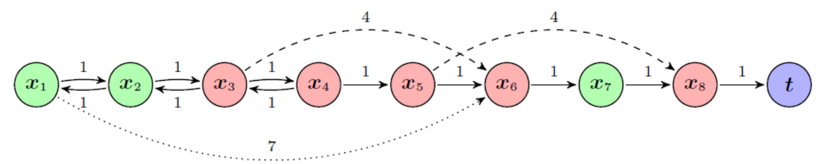

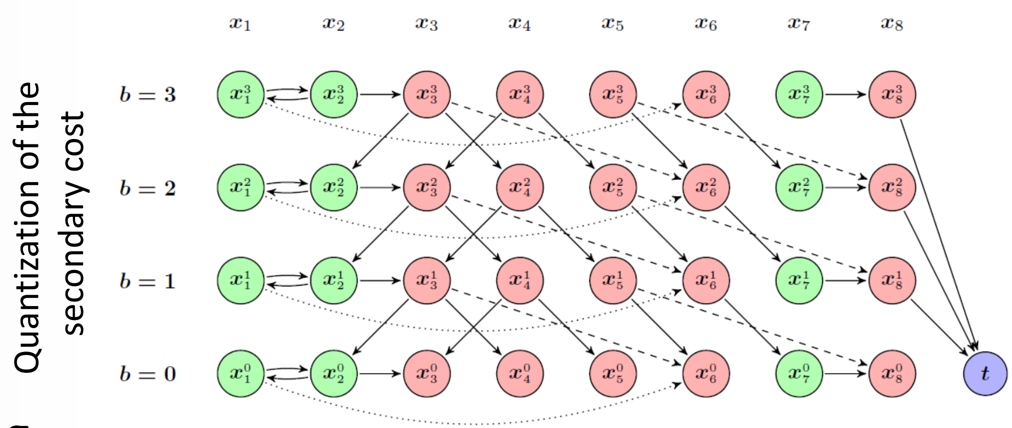

- One obtains the full Pareto curve
- For monotonic non-decreasing costs this graph can be search very efficiently

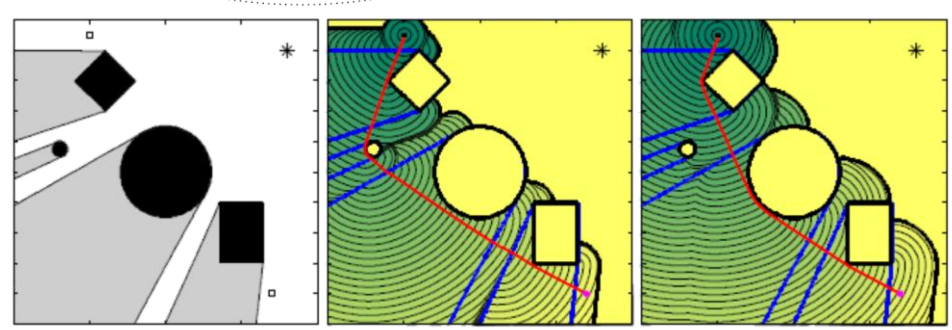

R. Takei, W. Chen, Z. Clawson, S. Kirov, and A. Vladimirsky, “Optimal control with budget constraints and resets”, SIAM Journal on Control and Optimization, to Appear.

Multi-Objective Planning Under Localization Constraints

- We are interested in a multi-objective problem where the secondary cost is a state dependent function
- In particular, taking into account strong priors, determine a path that minimizes length and position accuracy (never exceeding a maximum)

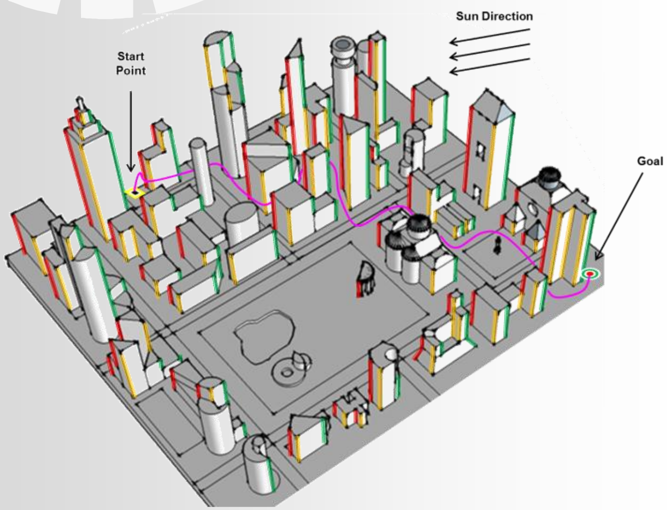

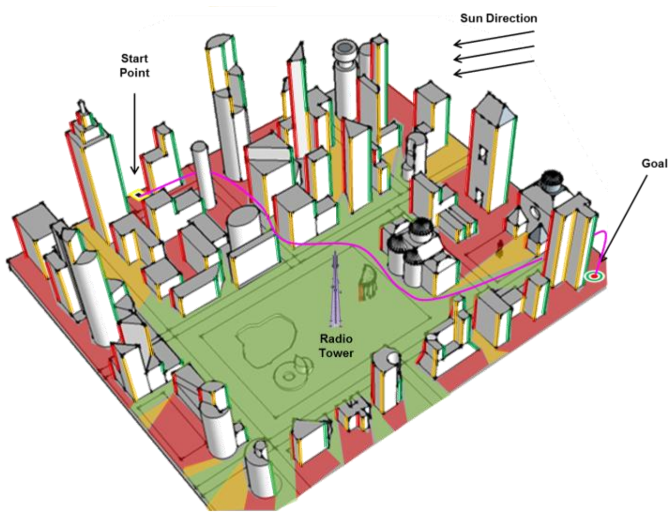

- Problem:

$$\min_{\pi_{sd} \in \mathcal{P}_{sd}} \quad C(\pi_{sd})$$

$$\text{s.t.} \quad \bar{\lambda}\Big(\underbrace{P(\pi_{sk})}_{\text{Cov. along path}} \Big) \leq p_{\max}\,, \forall \pi_{sk} \subseteq \pi_{sd} \in \mathcal{P}_{sd}\,,$$

Planning in Belief Space

This problem is related to work at MIT by Prof. Roy group

- Single objective:
  - Trace of the state estimate error covariance
  - Propagate the EKF over paths
  - Minimize uncertainty at the goal state

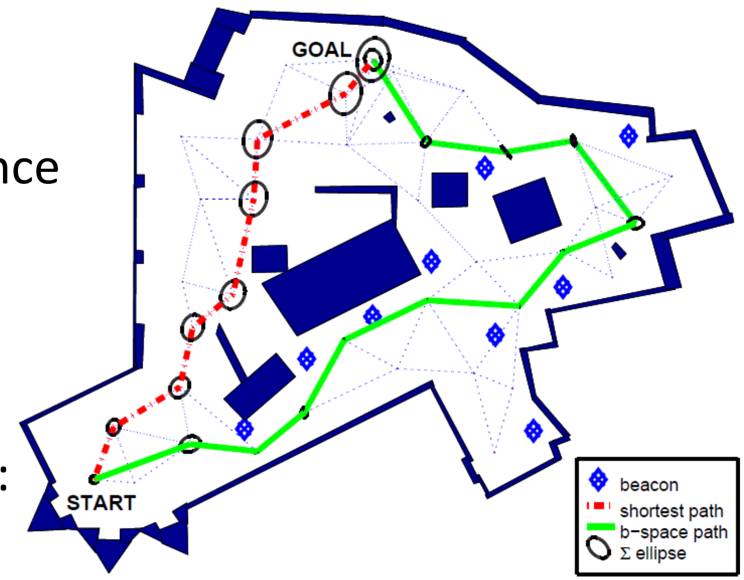

- Covariance factorization for fast computation:
- Write $P_t = B_t C_t^{-1}$ as

$$\begin{pmatrix} B_t \\ C_t \end{pmatrix} = \underbrace{\begin{pmatrix} F_t & Q_t F_t^{-T} \\ M(\gamma_t) F_t & F_t^{-T} + M(\gamma_t) Q_t F_t^{-T} \end{pmatrix}}_{W} \begin{pmatrix} B_{t-1} \\ C_{t-1} \end{pmatrix}$$

- Computation intensive as these weight matrix need be computed across the roadmap

Problem Setup

- We consider a general vehicle and sensing model

$$x(t+1) = f(x(t), n(t))$$
$$y_j(t) = h_j(x(t), v_j(t)), \qquad \forall j \in \{1, \dots, m(x,t)\},$$

- The error covariance for the Extended Kalman Filter

$$P_{t+1}^{-1} = (F_t P_t F_t' + Q_t)^{-1} + \sum_{j=1}^{m} H_j' R_j^{-1} H_j,$$

F, H : Jacobians, $\qquad Q, R$: Process/Measurement cov.

- We assume:
  - Data association is perfect and no misdetection
  - Consistency (mean state close to planned trajectory)

- To alleviate the computation burden of associate to each edge a matrix and propagate matrices over the edges we consider the *maximum eigenvalue* of the covariance matrix $\bar{\lambda}(P_t)$

Maximum Eigenvalue Bound

- Given a set of vertices in the roadmap
- Given a strong prior about the environment

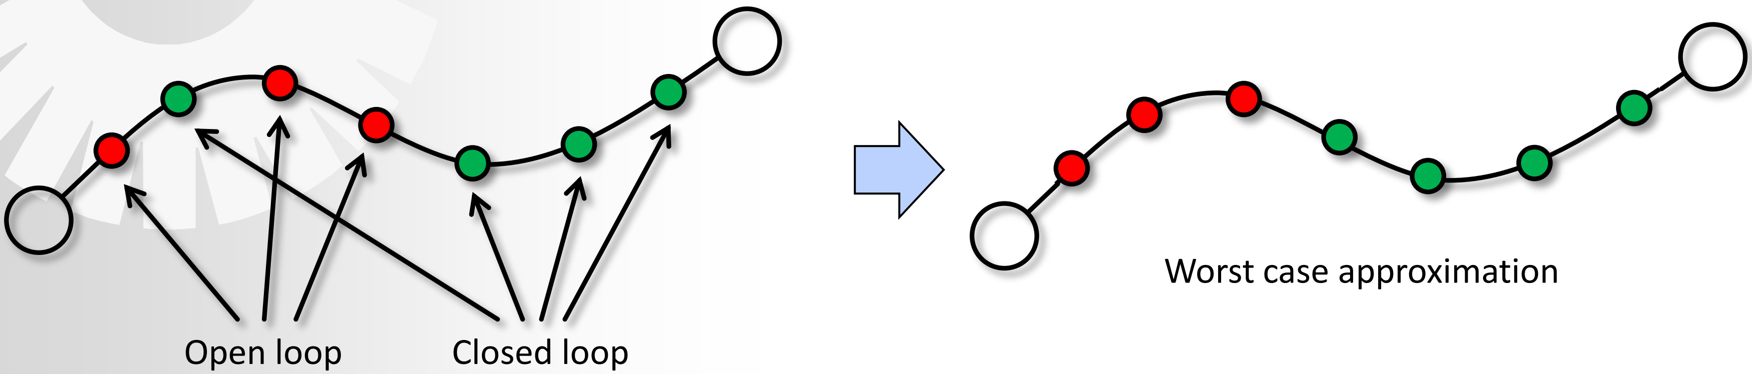

Theorem

$$\bar{\lambda}(P_T) \le b\kappa - \zeta + 1 \Big/ \left(\left(\frac{d-\zeta c}{\zeta c+1}\right)^{T-\kappa} \frac{1}{\zeta+\bar{\lambda}(P_0)} + \frac{c}{\zeta c+1}\left(\frac{1-\frac{(d-\zeta c)^{T-\kappa}}{(\zeta c+1)^{T-\kappa}}}{1-\frac{(d-\zeta c)}{(\zeta c+1)}}\right)\right)$$

where

$$b := \bar{\lambda}(Q)\,,\ c := \inf_{\hat{x}\in\mathcal{X}\setminus\mathcal{X}_S} \underline{\lambda}\left(H(f(\hat{x},0))' R^{-1} H(f(\hat{x},0))\right)$$

$$d := bc+1\,,\ \zeta := \left(bc+\sqrt{b^2c^2+4bc}\right)/(2c),\ \kappa := |\mathcal{X}_S|$$

Multi-Objective Planning with Localization Constraints

- The problem we are interested is the following:

$$\min_{\pi_{sd} \in \mathcal{P}_{sd}} \quad C(\pi_{sd})$$

$$\text{s.t.} \quad \bar{\lambda}(\underbrace{P(\pi_{sk})}_{\text{Cov. along path}}) \leq p_{\max}, \forall \pi_{sk} \subseteq \pi_{sd} \in \mathcal{P}_{sd},$$

- We can consider a similar approach as discussed before, i.e. solving the problem on an extended graph:

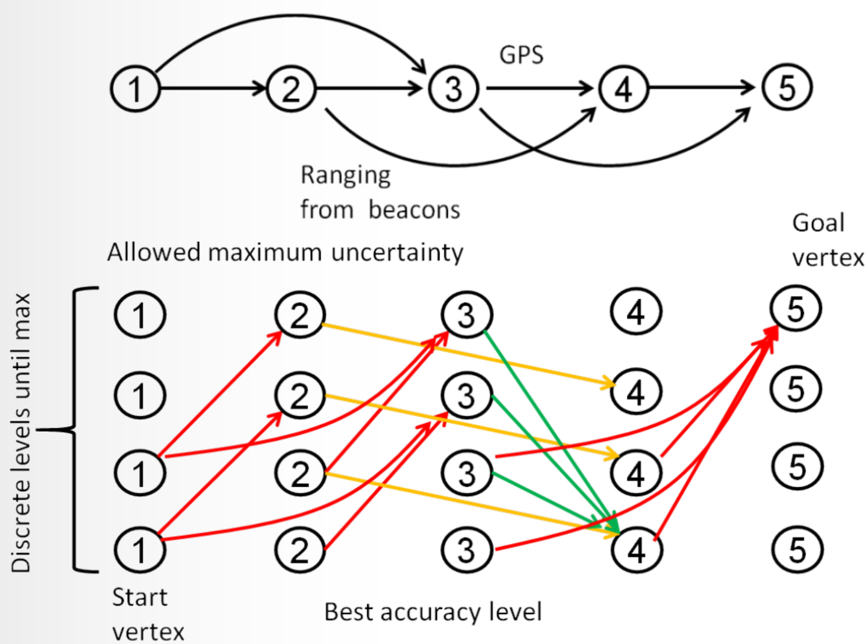

Simulations Results

- Sensor modalities: IMU + LIDAR to range to building corners

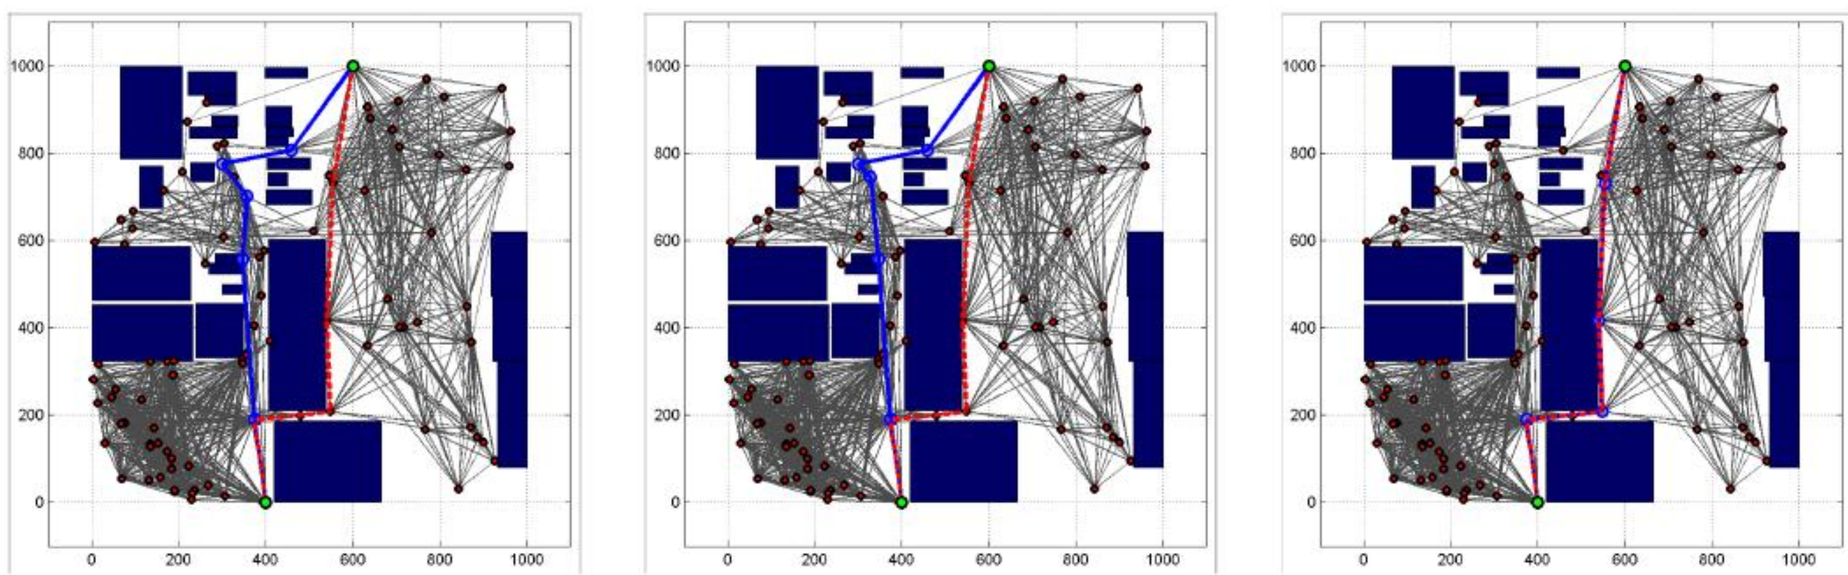

(a) $p_{\max} = 8$ (b) $p_{\max} = 9$ (c) $p_{\max} = 13$

The extended graph can become very large

- Planning in a $1km^2$ environment
- 100 vertices on the PRM
- ~2000 edges

$p_{\max}$	Uniform	
	Edges	Nodes
8	174542	9257
9	205268	10486
12	297642	14207
13	328368	15436
14	359192	16682

How does one choose the quantization level for the secondary cost?

Sparse Extended Graph

- Consider the change of $\bar{\lambda}(P_0)$ over an edge $e \in \mathcal{R}$

$$\Delta_e(\bar{\lambda}(\mathbf{P}_0)) := B_e(\bar{\lambda}(\mathbf{P}_0)) - \bar{\lambda}(\mathbf{P}_0)$$

then for each edge e we can compute the worst-case difference Δ_e^* as this function is concave

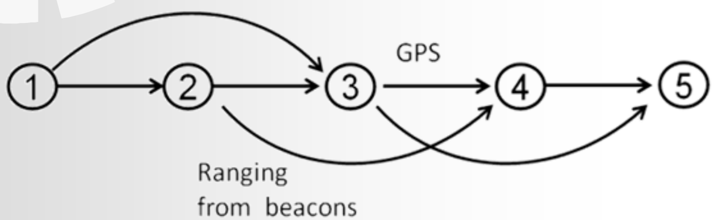

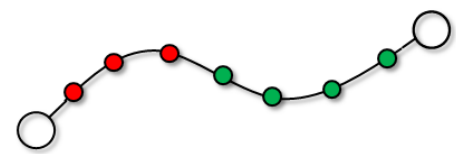

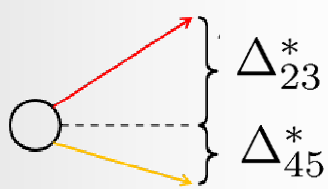

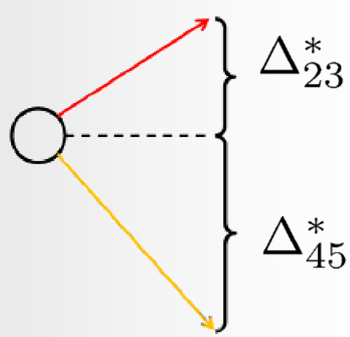

$$\delta := \min_{e \in E}\{|\Delta_e^*| \; : \; |\Delta_e^*| > 0\}$$

Two Schemes

- Uniform $\delta := \min_{e \in E} \{ |\Delta_e^*| : |\Delta_e^*| > 0 \}$

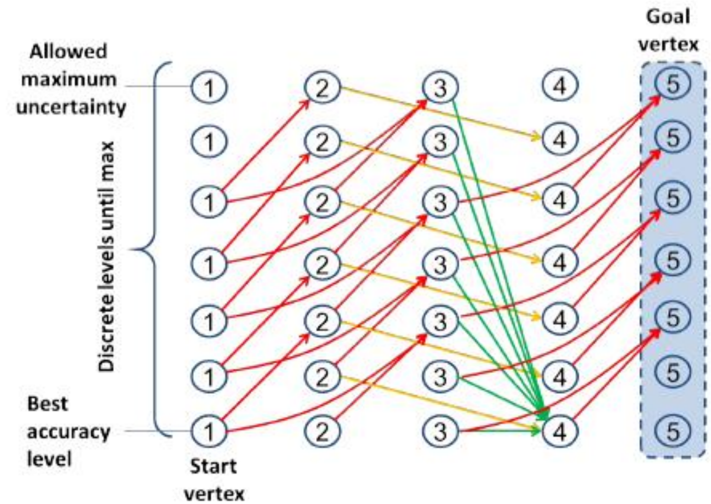

- Adaptive $\delta_j := \min_{i \in \mathcal{N}_{in}(j)} \{ |\Delta_{e_{ij}}^*| : |\Delta_{e_{ij}}^*| > 0 \}$

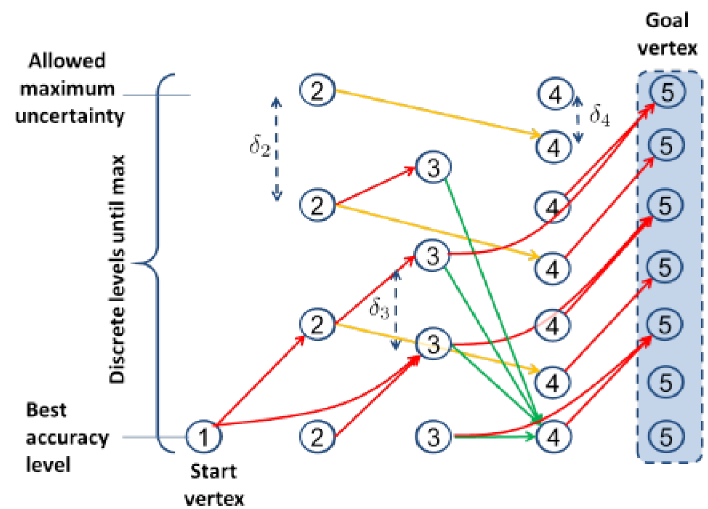

Results for Adaptive Scheme

- Sensor modalities: IMU + LIDAR to range to building corners

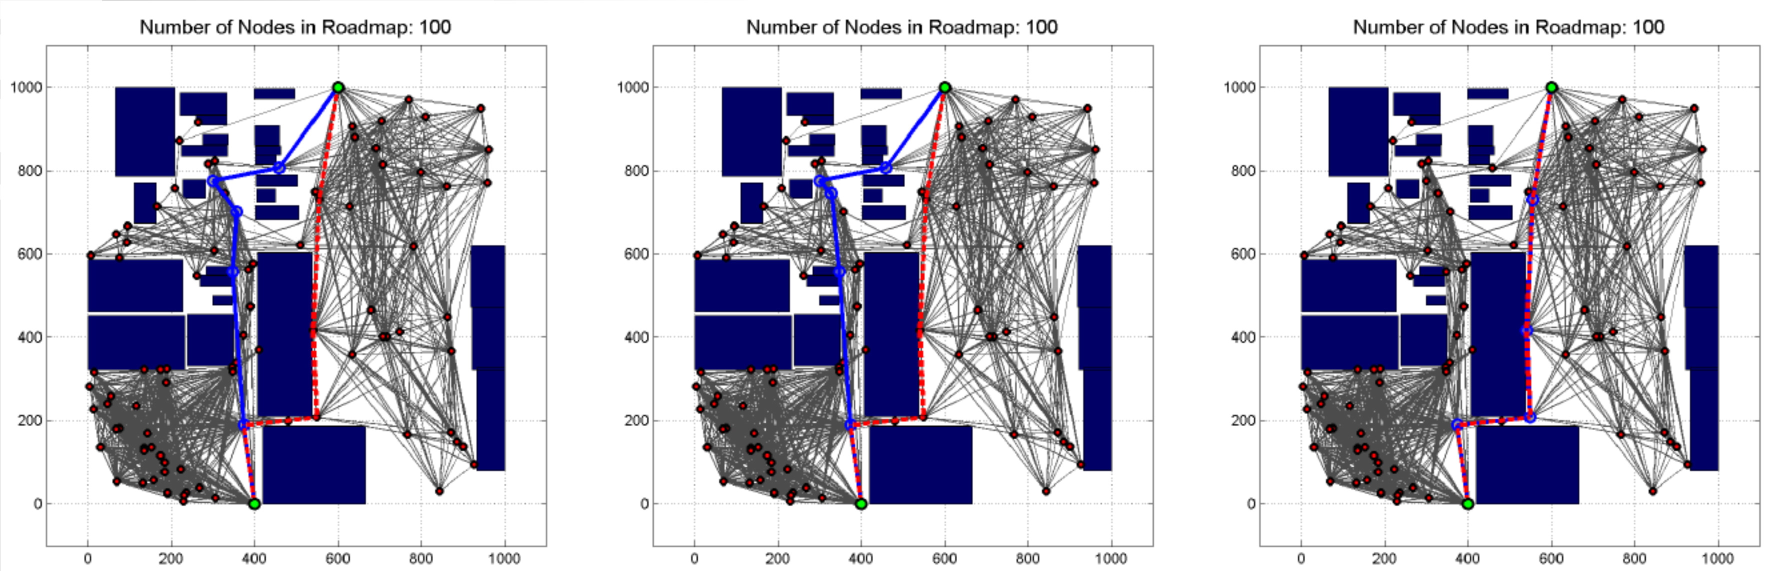

(a) $p_{\max} = 8$ (b) $p_{\max} = 9$ (c) $p_{\max} = 13$

The extended graph can become very large

- Planning in a $1km^2$ environment
- 100 vertices on the PRM
- ~2000 edges

$p_{\max}$	Adaptive		Uniform		Improvement
	Edges	Nodes	Edges	Nodes	
8	47413	2253	174542	9257	3.68
9	55354	2541	205268	10486	3.71
12	79104	3405	297642	14207	3.76
13	87023	3693	328368	15436	3.77
14	94930	3970	359192	16682	3.78

Interactions

- Recall the hierarchical planning

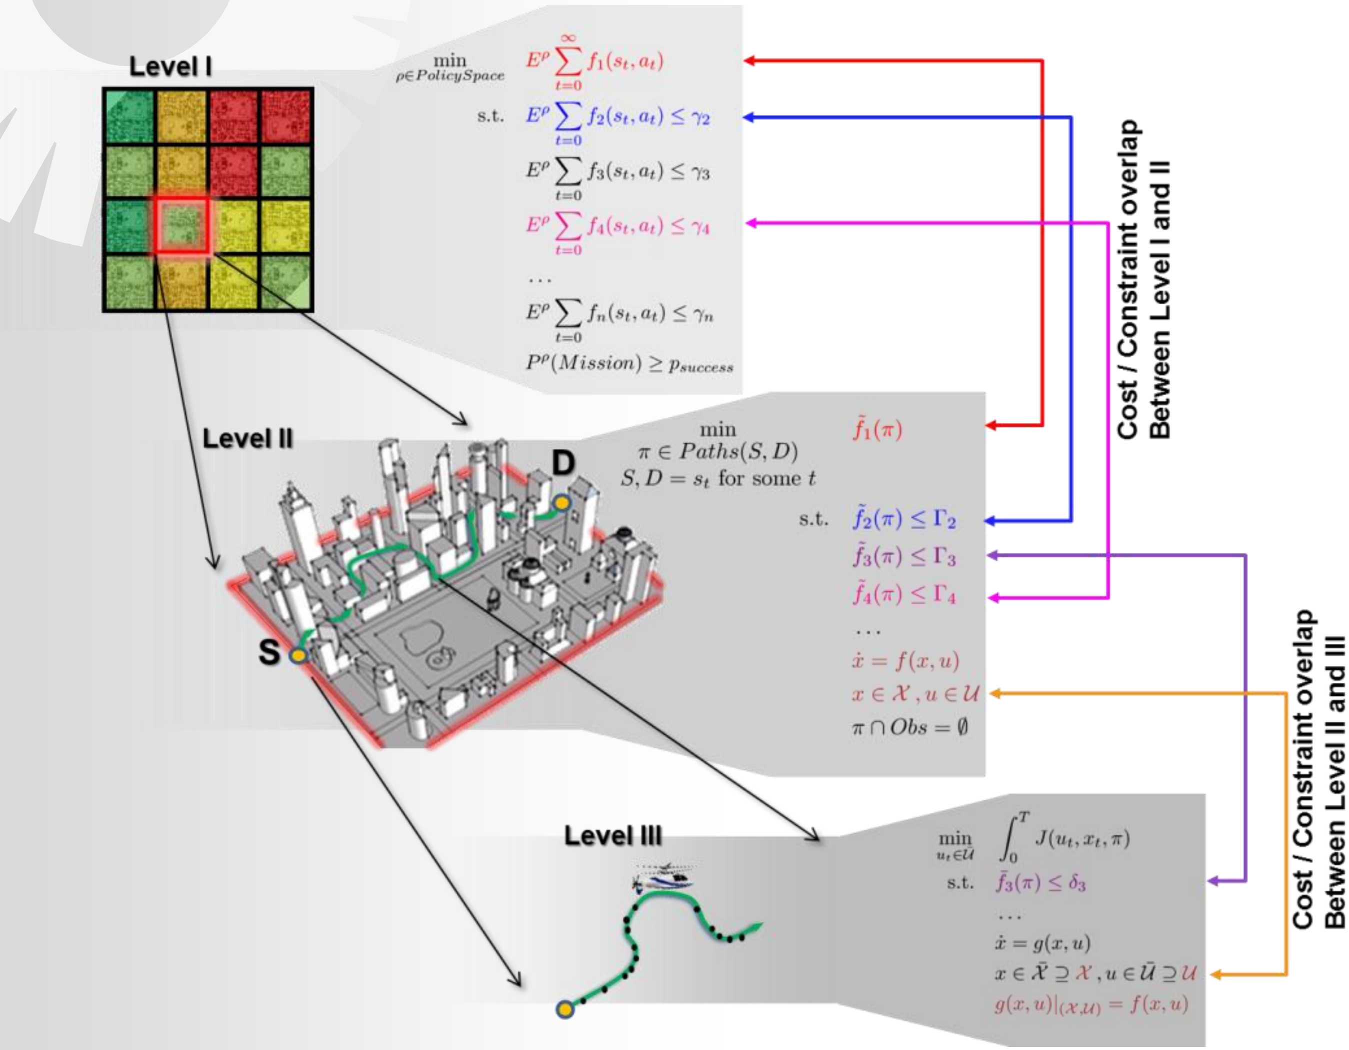

Example: Interaction Between Mission and Motion Planning

1. Mission planning determine optima policy to have autonomous system go from Start to Goal with constraints on missions success and threat exposure
2. When new threats are found, interaction between planners lead to a new mission level policy

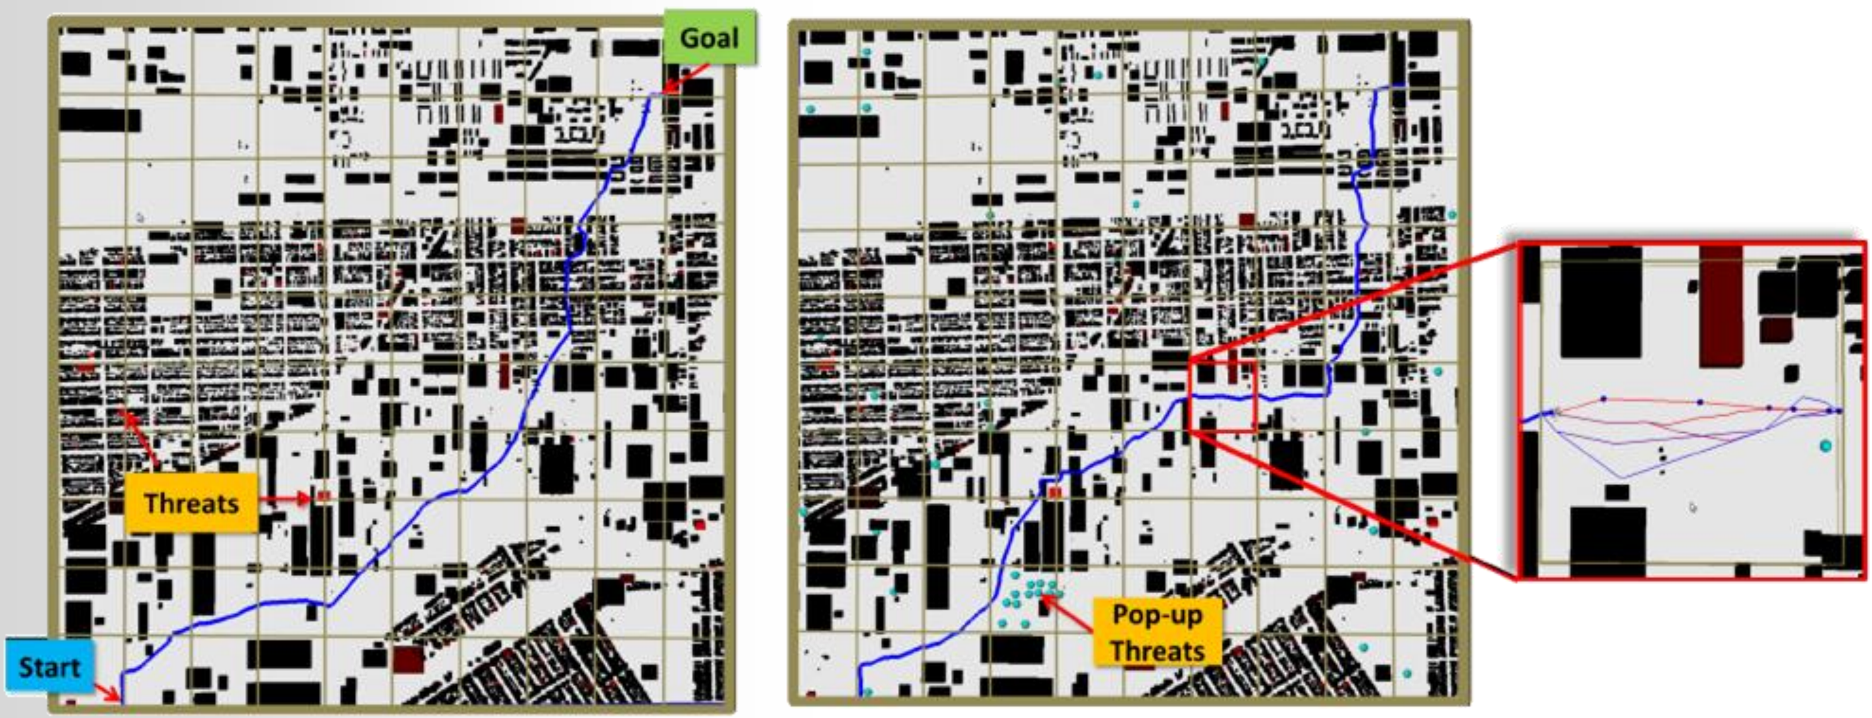

Conclusions

- Hierarchical Planning
  - Mission planning from LTL specifications define a policy at coarse scale
  - Motion planning enables navigation in complex environments
  - "Coupled" multi-objective planning algorithms enable autonomous vehicle to deal with contingencies at multiple temporal and spatial scales

- Multi-objective path planning
  - Developed a new algorithms that find a path in a complex environment that minimizes multiple costs
  - Explored computation/accuracy tradeoffs to ensure algorithms can be implemented in real-time.

For more information contact:

Alberto Speranzon

Research Scientist
Control Systems Group
Systems Department

sperana@utrc.utc.com