An Experimental Study of Robust Distributed Multi-Robot Data Association from Arbitrary Poses

Erik Nelson[1] **Vadim Indelman**[2] **Nathan Michael**[1] **Frank Dellaert**[2]

[1]Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15232

[2]College of Computing
Georgia Institute of Technology
Atlanta, GA 30332

June 17, 2014

Motivating Scenario

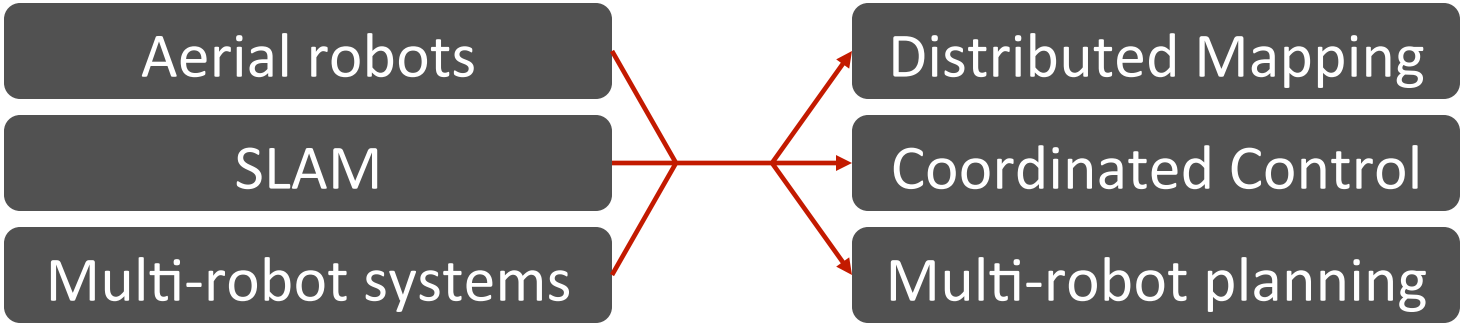

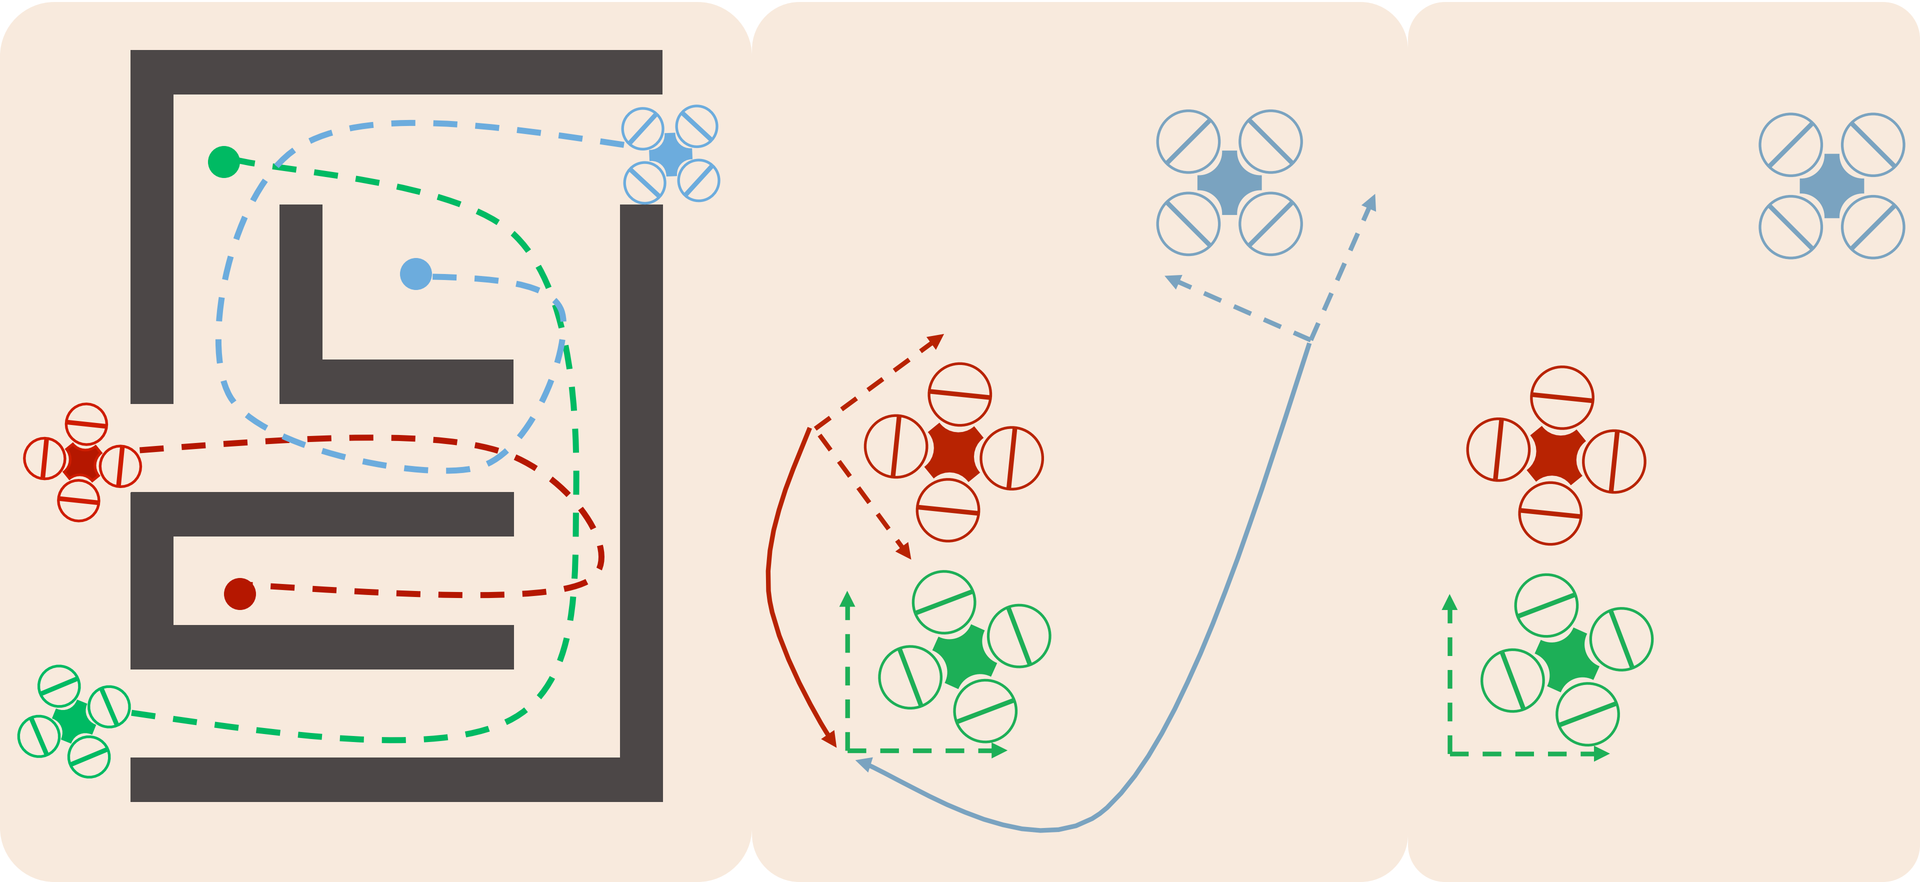

This Work

Experimentally evaluates and extends [Indelman, et al., ICRA 2014]

1. Related and prior work
2. Technical approach from [Indelman, et al., ICRA 2014]
3. Algorithmic complexity, metrics for saliency of information
4. Experimental design
5. Transformation accuracy experiments
6. Network complexity and run time efficiency experiments

Related and Prior Work

Landmark based

Landmarks and waypoints observed throughout an environment localize each robot to the same coordinate frame

[Fenwick, et al., ICRA 2002]
[Olson, et al., IEEE J. Oceanic Engineering 2006]

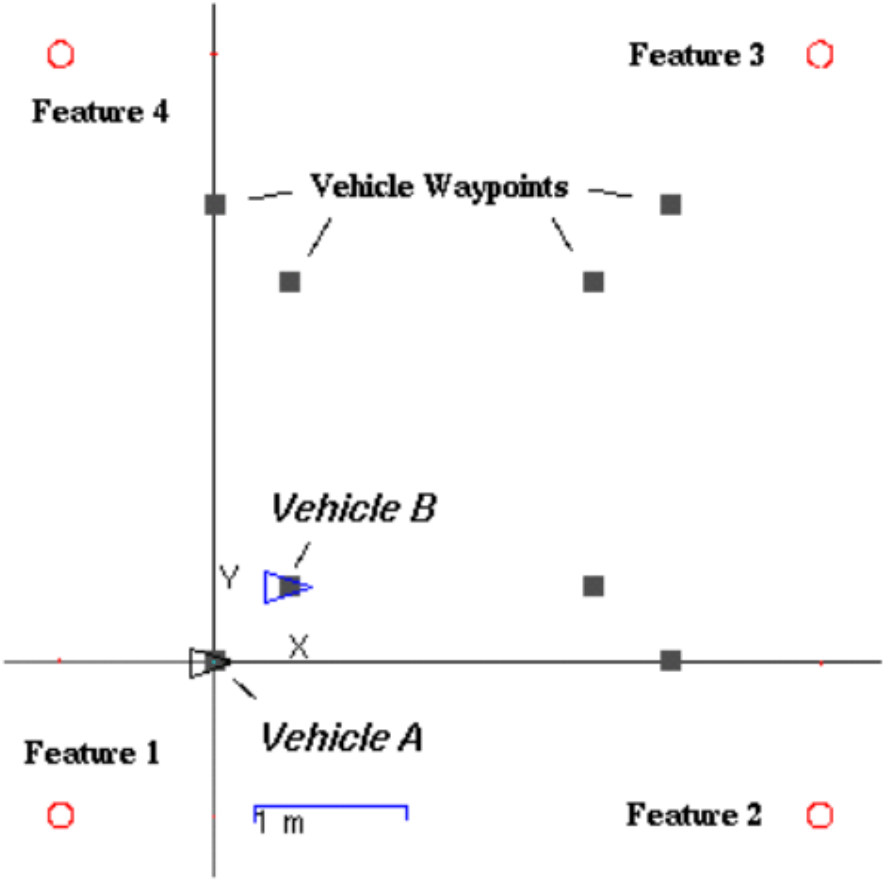

Fenwick, et al., ICRA 2002

Related and Prior Work

Landmark based

Landmarks and waypoints observed throughout an environment localize each robot to the same coordinate frame

[Fenwick, et al., ICRA 2002]
[Olson, et al., IEEE J. Oceanic Engineering 2006]

Direct inter-robot observations

Robots observe one another

[Kim, et al., ICRA 2010]
[Bailey, et al., ICRA 2011]
[Howard, et al., IJRR 2006]
[Zhou, et al., IROS 2006]
[Charrow, et al., ISER 2012]

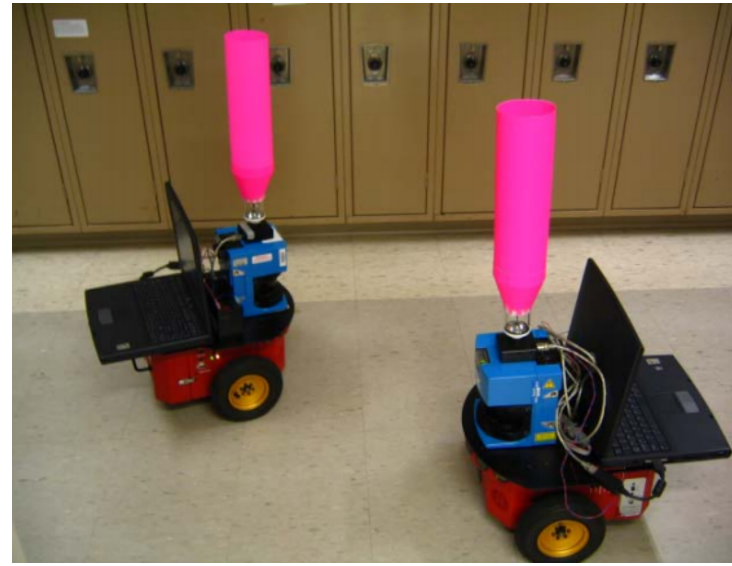

Zhou, et al., IROS 2006

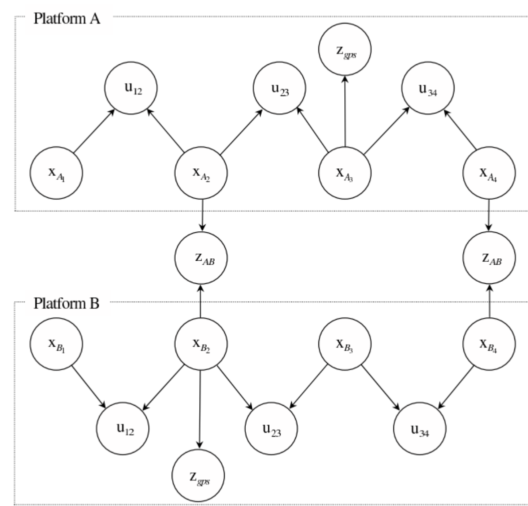

Bailey, et al., ICRA 2011

Related and Prior Work

Landmark based

Landmarks and waypoints observed throughout an environment localize each robot to the same coordinate frame

[Fenwick, et al., ICRA 2002]
[Olson, et al., IEEE J. Oceanic Engineering 2006]

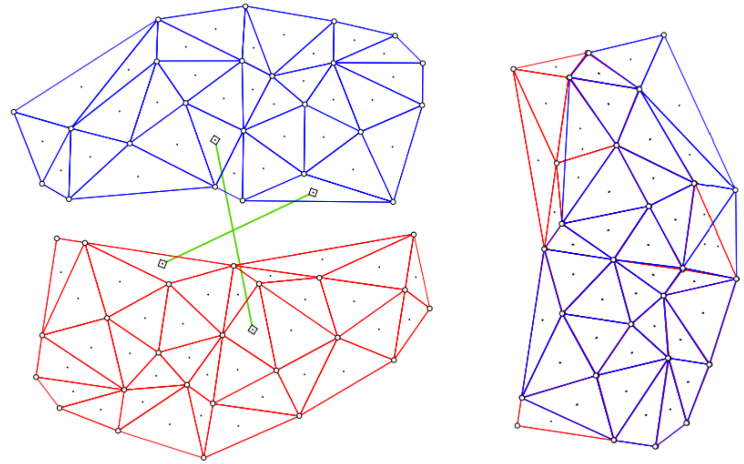

Cunningham, et al., ICRA 2012

Direct inter-robot observations

Robots observe one another

[Kim, et al., ICRA 2010]
[Bailey, et al., ICRA 2011]
[Howard, et al., IJRR 2006]
[Zhou, et al., IROS 2006]
[Charrow, et al., ISER 2012]

Data association

Localization using correspondences formed between data shared by robots

[Montijano, et al., IEEE Trans. Robotics, 2013]
[Cunningham, et al., ICRA 2012]
[Indelman, et al., ICRA 2014]

Technical Approach

[Indelman, et al., ICRA 2014]

Goal:

- Establish multi-robot data association
- Infer initial relative poses

Strategy:

- Robots share observations
- Calculate candidate multi-robot relative pose constraints
- Collect into set, $\mathscr{F}$, of correspondences (includes many outliers)
- Use EM to estimate inlier correspondences while inferring relative initial poses for each robot

Local trajectories of 3 robots

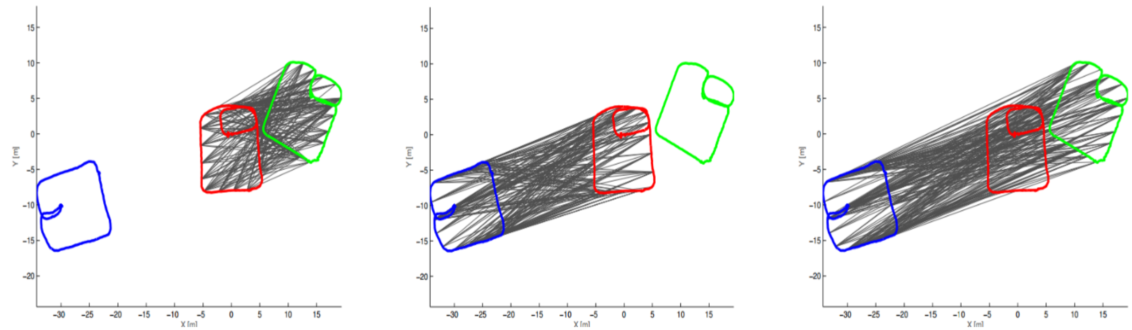

Estimated

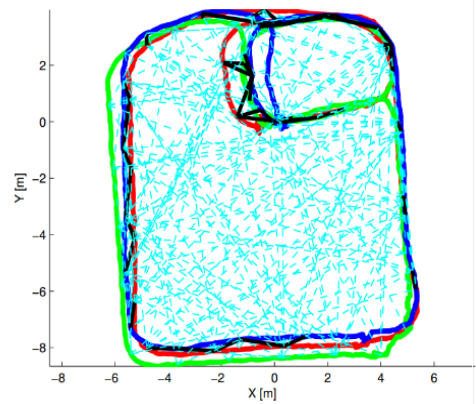

Ground truth

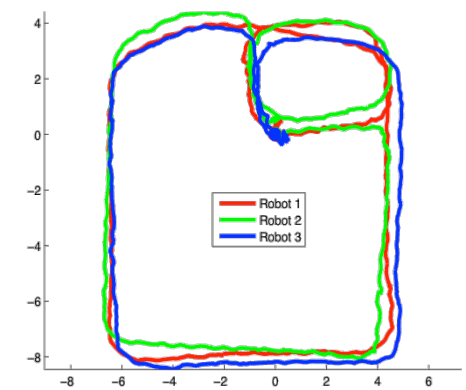

Technical Approach

Multi-robot system represented as a factor graph

Data associations, $(r_i, r_j, k, l) \in \mathscr{F}$, represent pose constraints, $c_{k,l}^{r_i,r_j}$

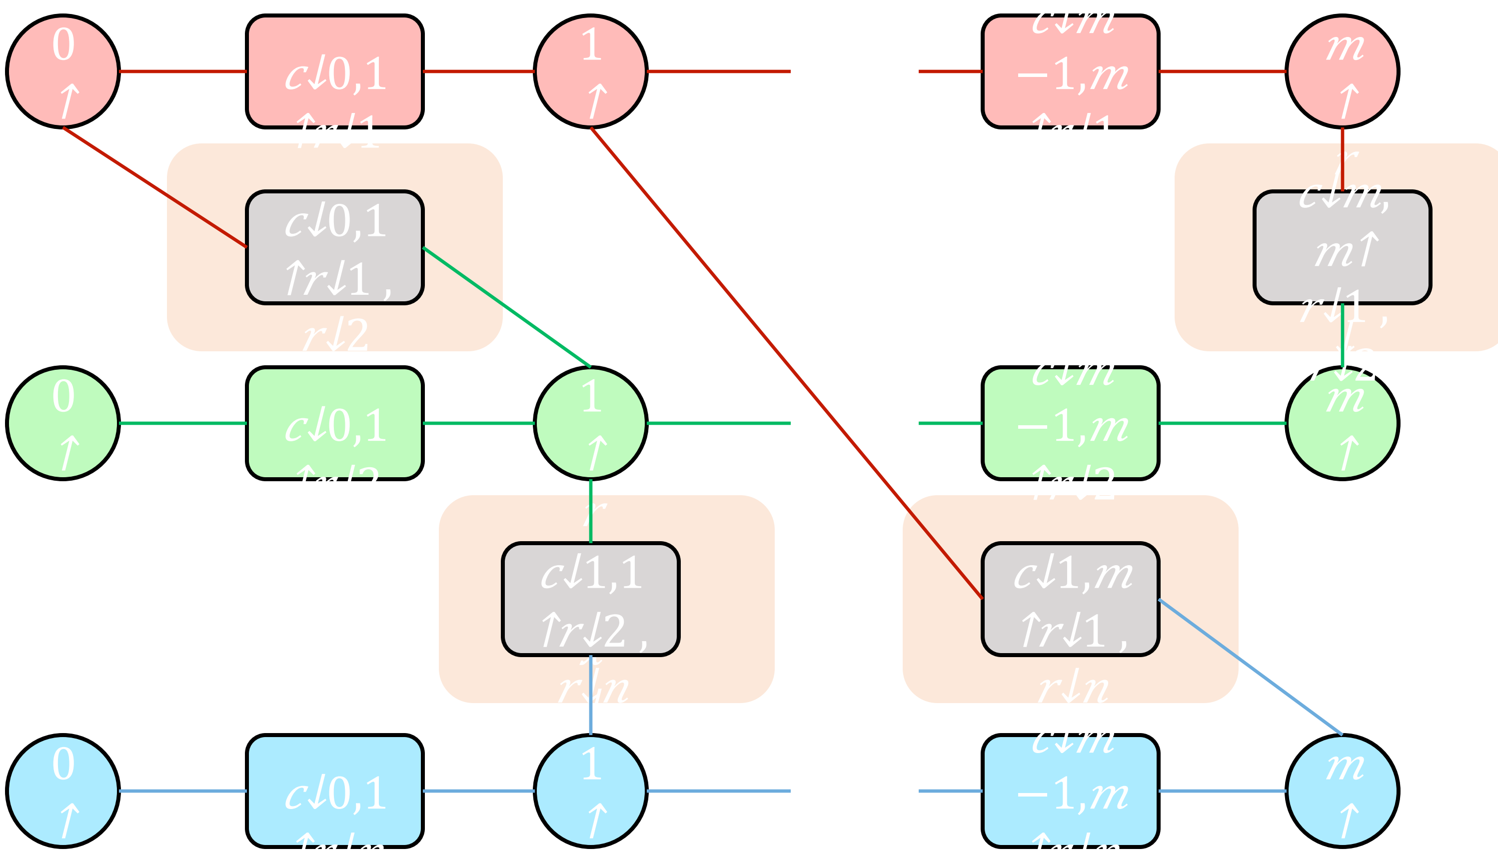

Technical Approach

Multi-robot joint pdf:

$$p\left(X|Z\right) \propto \prod_r p\left(X^r|Z^r\right) \prod_{(r_i,r_j,k,l)\in\mathscr{F}} p\left(c_{k,l}^{r_i,r_j} | x_k^{r_i}, x_l^{r_j}\right)$$

Local measurements

Data association

- X^r - trajectory of robot r
- Z^r - robot r's observations
- $\mathscr{F}$ - set of data associations

Multi-robot measurement likelihood

$$p\left(c_{k,l}^{r_i,r_j} | x_k^{r_i}, x_l^{r_j}\right) \propto \exp\left(-\frac{1}{2}\left\| err\left(c_{k,l}^{r_i,r_j}, x_k^{r_i}, x_l^{r_j}\right)\right\|_\Sigma^2\right)$$

$$err\left(c_{k,l}^{r_i,r_j}, x_k^{r_i}, x_l^{r_j}\right) \doteq c_{k,l}^{r_i,r_j} \ominus h\left(x_k^{r_i}, x_l^{r_j}\right)$$

$$h\left(x_k^{r_i}, x_l^{r_j}\right) \doteq x_k^{r_i} \ominus \left(\boxed{T_{r_j}^{r_i}} \oplus x_l^{r_j}\right)$$

Unknown

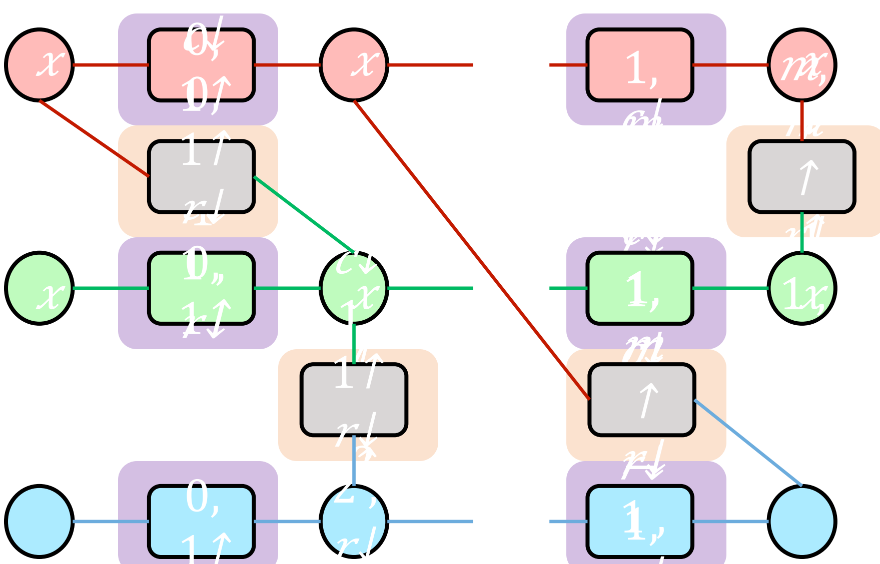

$a \ominus b$ – subtraction with b expressed in the frame of a

$\oplus$ – transformation composition

Technical Approach

Consider a robot pair

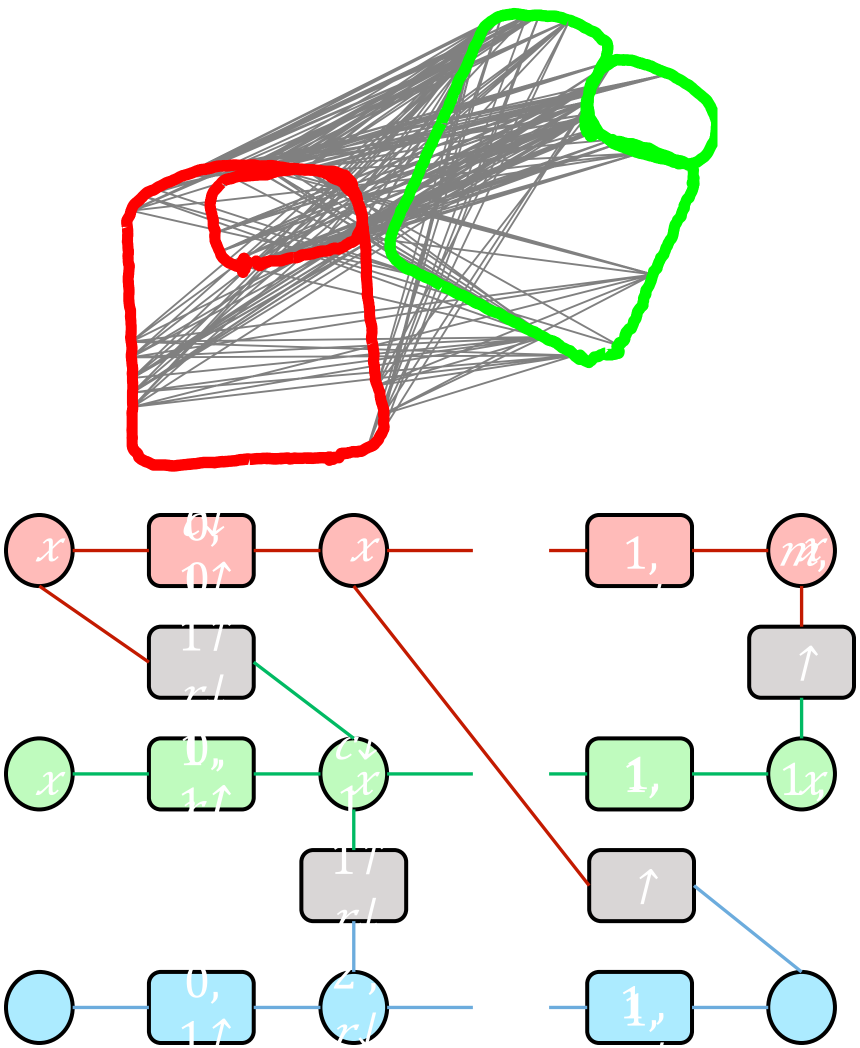

Initial relative pose estimates

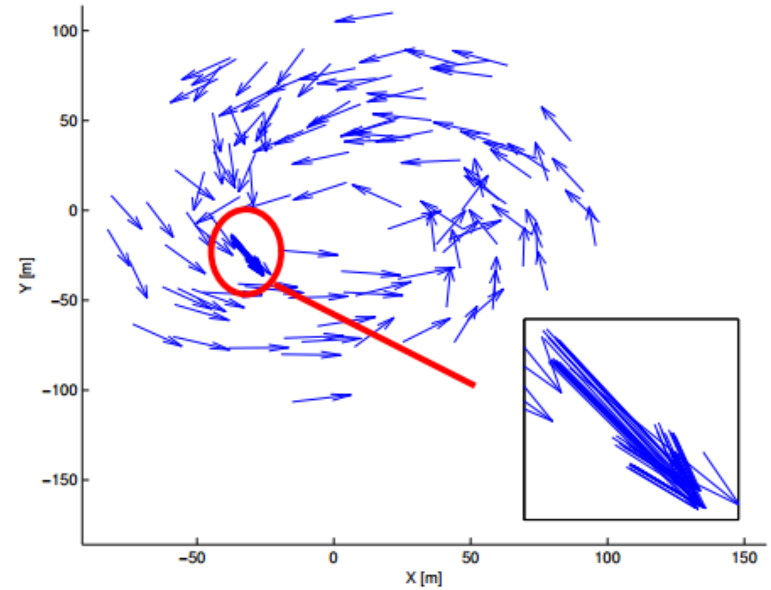

Relative initial pose estimates can be estimated from each candidate multi-robot correspondence

But only inliers yield similar transformations

E: estimate inlier correspondences given $T_{r_j}^{r_i}$

M: maximize over $T_{r_j}^{r_i}$ given inlier estimates to update $T_{r_j}^{r_i}$

Complexity and Saliency of Information

Problem: Run time complexity of sharing observations is $O(n^2 m^2)$

- n – robots
- m – shared observations per robot

Hypothesis: Selecting only the most salient observations will mildly reduce transformation accuracy while drastically increasing efficiency.

?

Algorithm LaserAutocovariance($z, iters$)

$reference \leftarrow z.$

$\Sigma \leftarrow 0_{3\times 3}.$

for $i = 1$ to $iters$ **do**

$\quad target \leftarrow \mathrm{Perturb}\,(reference, \sigma_x, \sigma_y, \sigma_\theta).$

$\quad \Sigma \leftarrow \Sigma + \mathrm{ICP\_Covariance}\,(target, reference).$

end

return $(\mathrm{Trace}\,(\Sigma))^{-1}$

[Nieto et al., Robot. Auton. Syst., 2007]

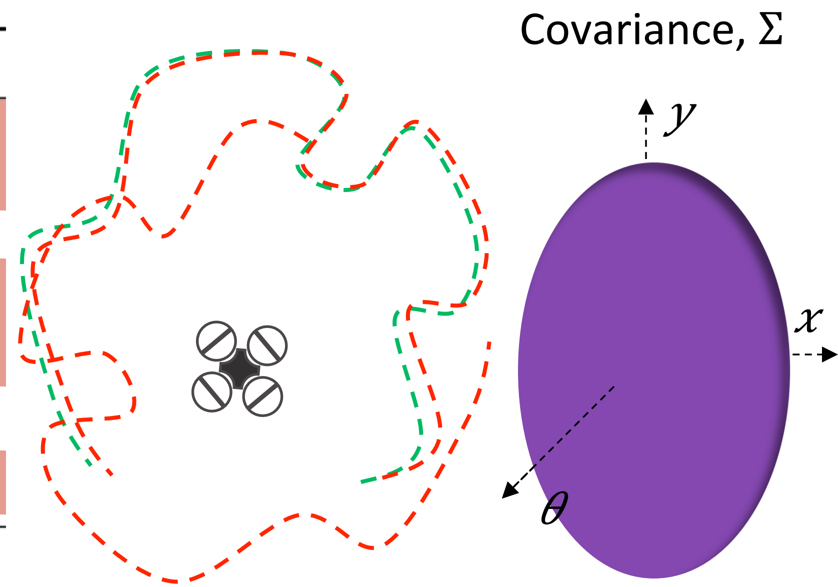

Complexity and Saliency of Information

Laser scan saliency, computed via autocovariance

Locations with high numbers of ICP correspondences

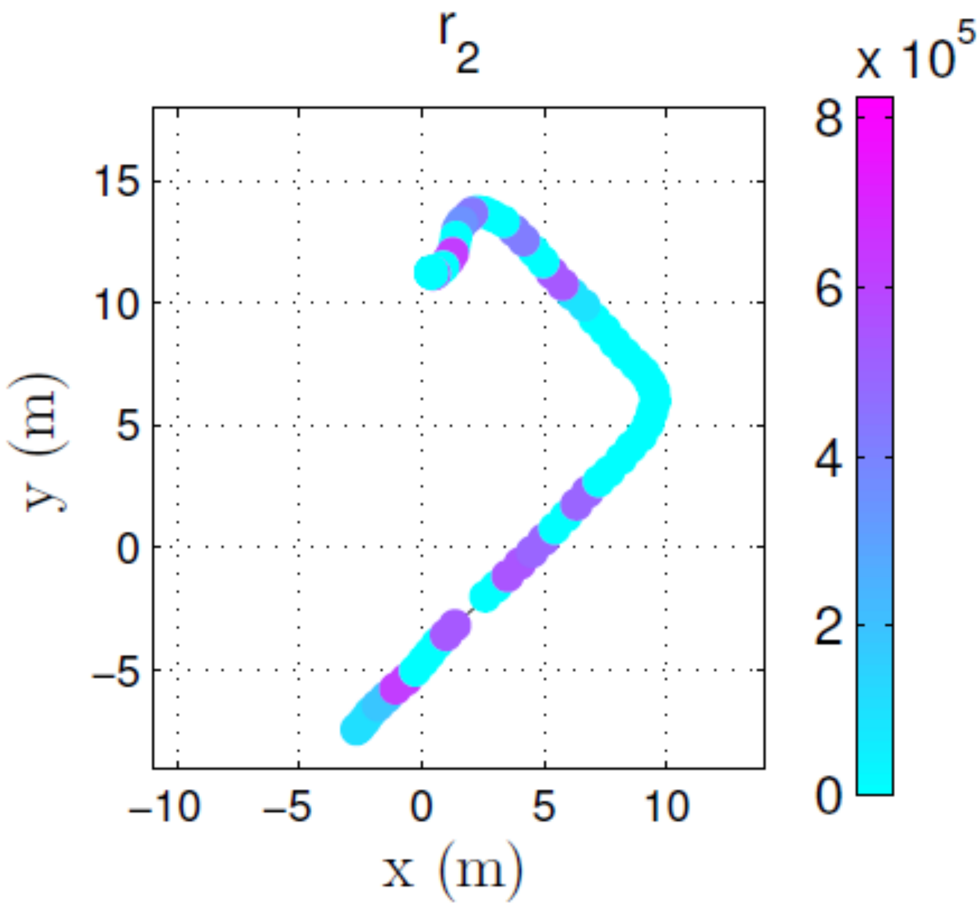

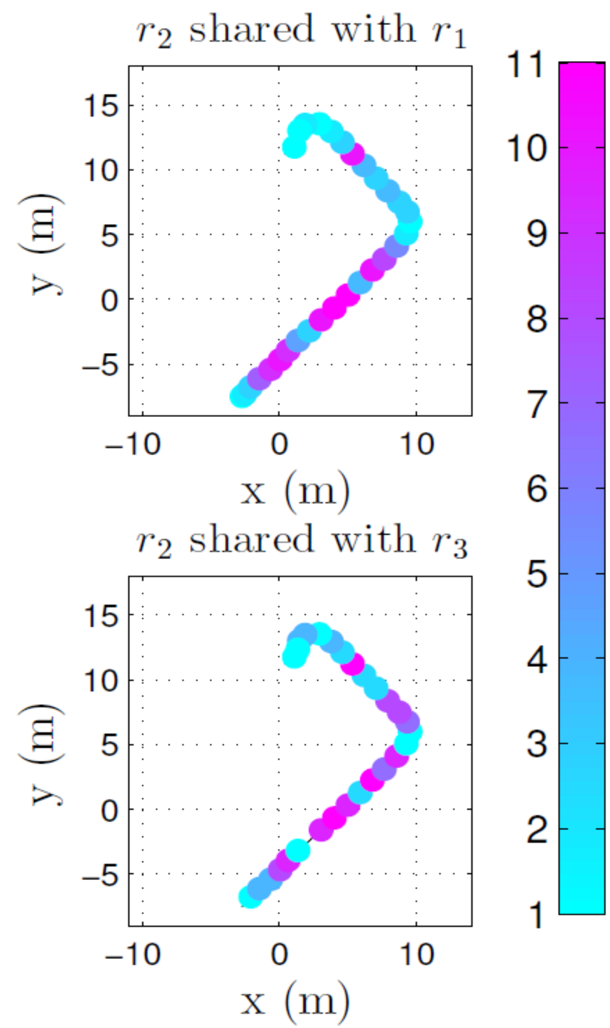

Complexity and Saliency of Information

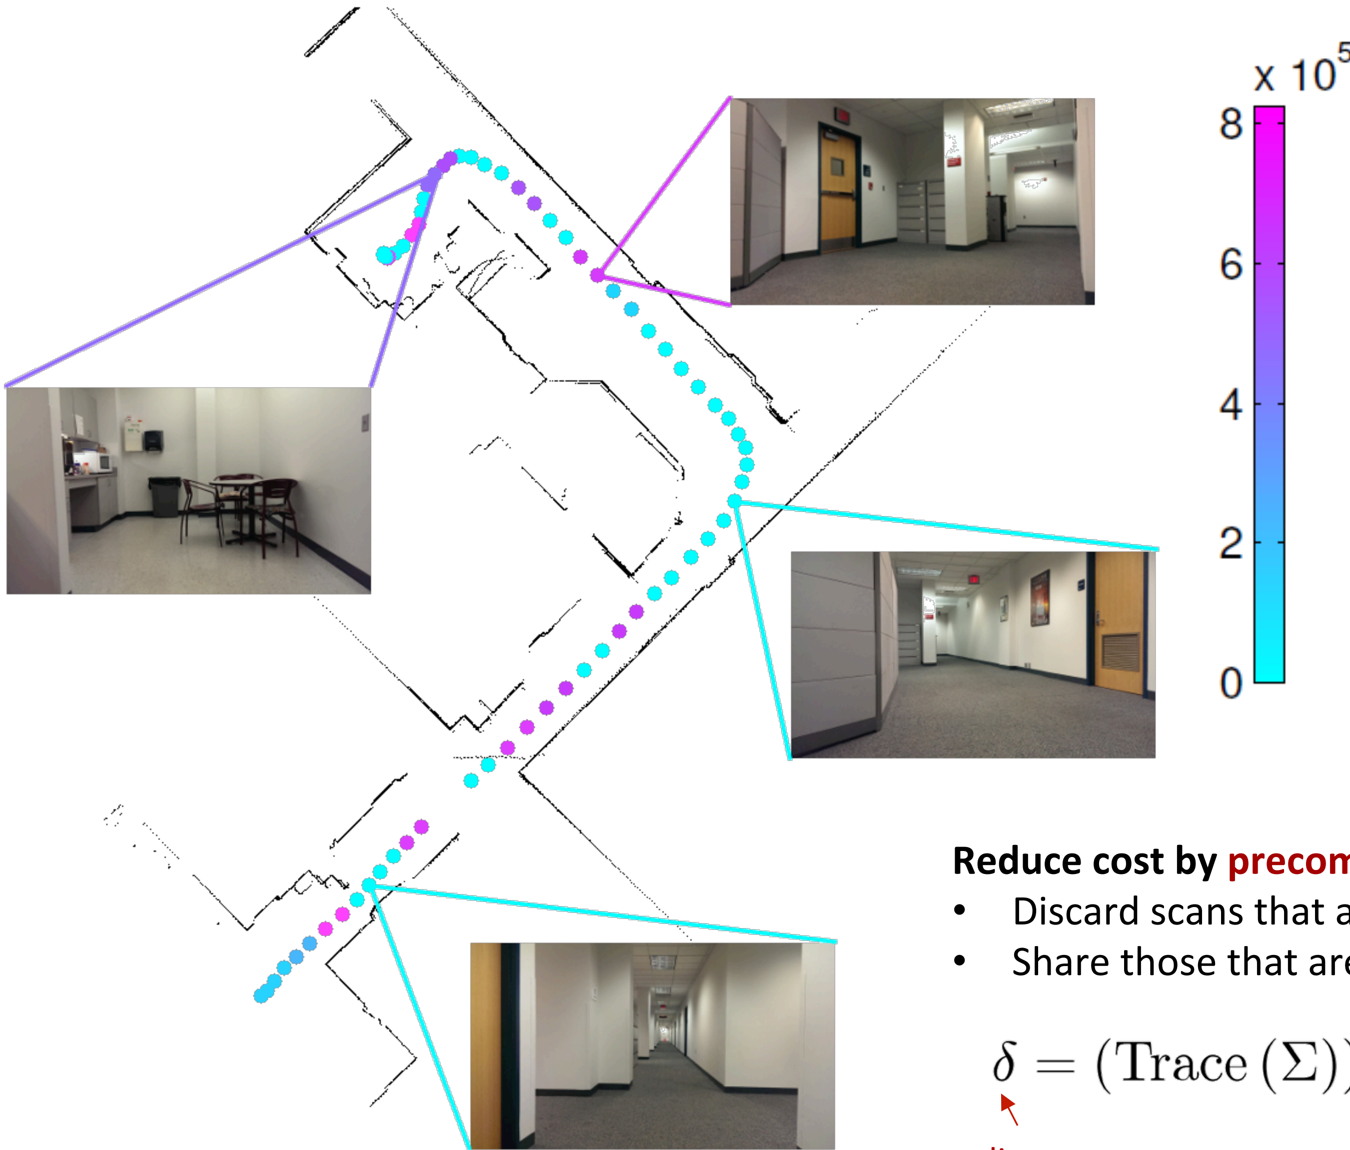

Reduce cost by precomputing observation saliency

- Discard scans that aren't salient
- Share those that are

$$\delta = (\mathrm{Trace}\,(\Sigma))^{-1}, \text{share if } \delta > \delta_s$$

saliency

threshold

Experimental Design and Approach

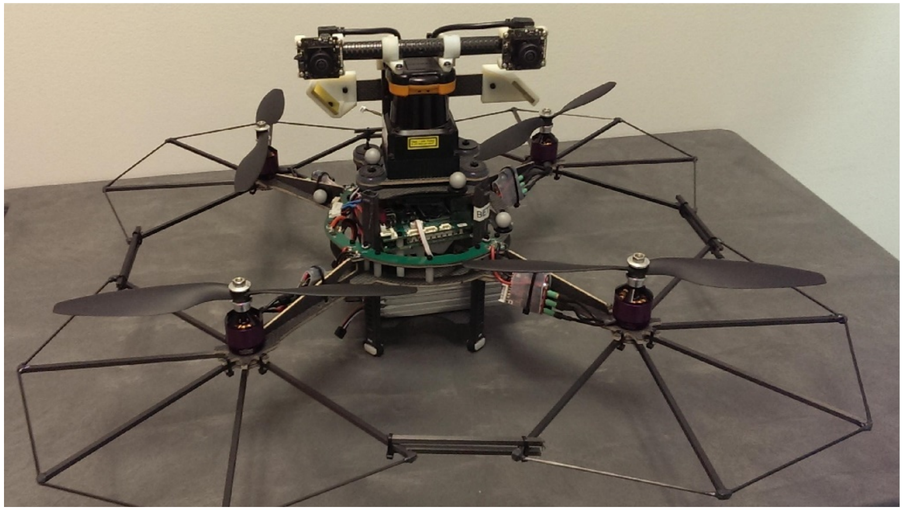

Platform

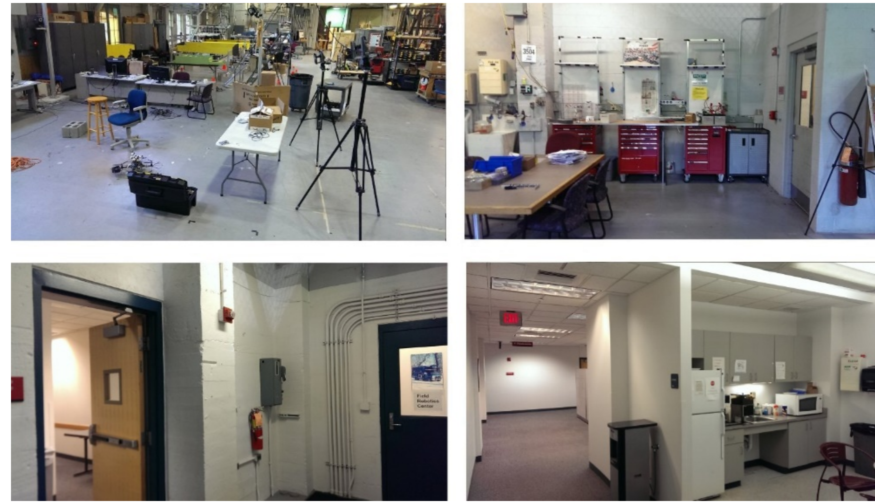

Trial **T1**

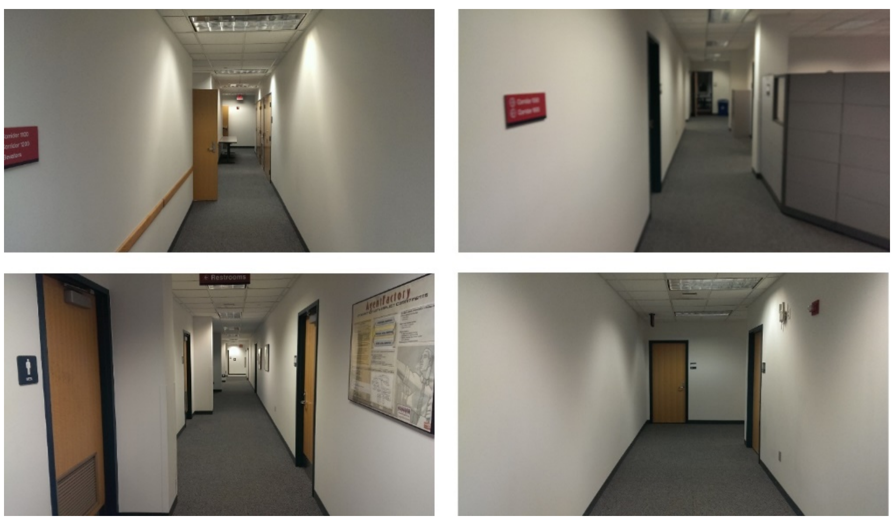

Trial **T2**

Trial **T3**

Experimental Design and Approach

SLAM implementation with a single robot

Results: Transformation Accuracy

T1 robot trajectories aligned with computed transformations

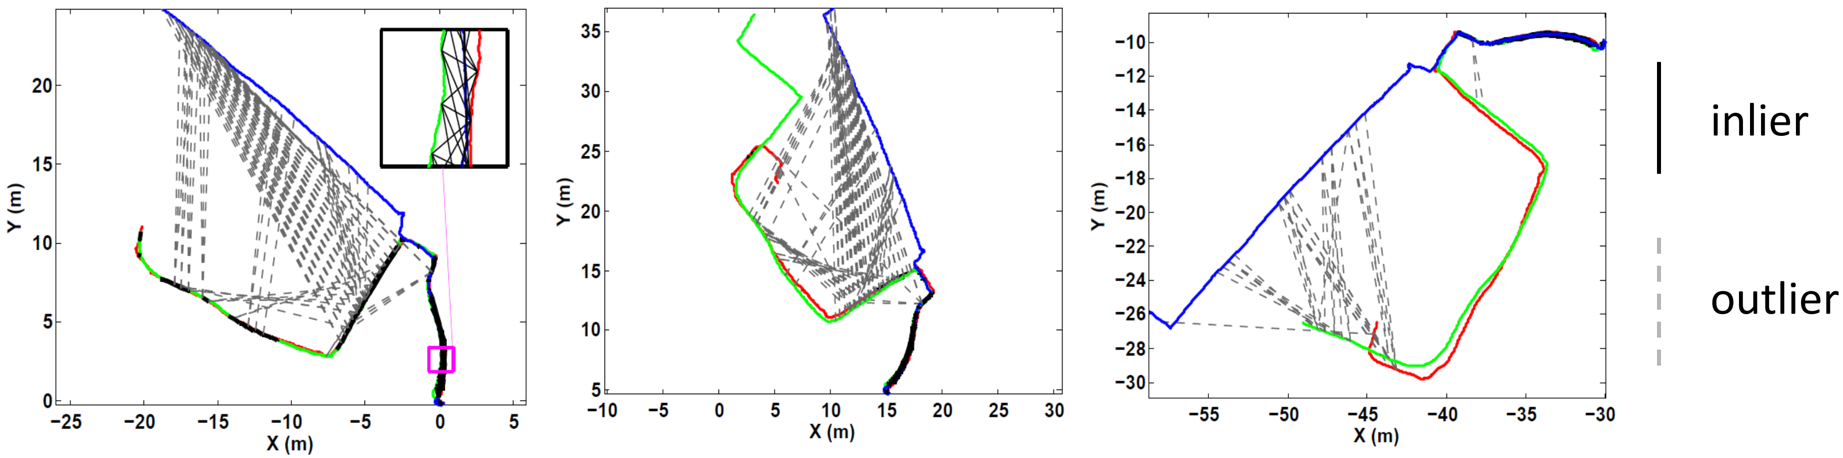

r_1 local frame (red) r_2 local frame (green) r_3 local frame (blue)

Computed and measured transformations

			Trial **T1**		Trial **T2**		Trial **T3**	
			$T^{r_1}_{r_2}$	$T^{r_1}_{r_3}$	$T^{r_1}_{r_2}$	$T^{r_1}_{r_3}$	$T^{r_1}_{r_2}$	$T^{r_1}_{r_3}$
Computed	x	(m)	-0.12	0.15	2.62	-4.53	1.41	-13.59
	y	(m)	-0.03	-0.27	7.45	-4.09	-3.99	-1.24
	θ	(rad)	-0.02	0.03	-1.57	0.00	0.97	2.05
Measured	x	(m)	0.00	0.00	2.48	-4.60	1.42	-13.63
	y	(m)	0.00	0.00	7.50	-3.99	-3.90	-1.02
	θ	(rad)	0.00	0.00	-1.57	0.00	1.08	2.01
Error	$\|\|x,y\|\|$	(m)	0.12	0.31	0.15	0.12	0.09	0.22
	θ	(rad)	0.02	0.03	0.00	0.00	0.11	0.04

Results: Saliency Thresholding

Computed and measured transformation errors

$\sim$ - No transformation established

	Trial **T2**							Trial **T3**						
δ_s	Shared scans			$T_{r_2}^{r_1}$ error		$T_{r_3}^{r_1}$ error		Shared scans			$T_{r_2}^{r_1}$ error		$T_{r_3}^{r_1}$ error	
	r_1	r_2	r_3	$\lVert x,y \rVert$ (m)	θ (rad)	$\lVert x,y \rVert$ (m)	θ (rad)	r_1	r_2	r_3	$\lVert x,y \rVert$ (m)	θ (rad)	$\lVert x,y \rVert$ (m)	θ (rad)
0	75	77	65	0.15	0.00	0.20	0.00	74	55	71	0.09	0.10	0.22	0.05
2×10^5	22	26	23	0.19	0.00	0.24	0.00	26	18	36	0.22	0.08	0.59	0.13
4×10^5	22	24	23	0.19	0.00	0.24	0.00	24	16	35	$\sim$	$\sim$	0.59	0.13
6×10^5	16	18	19	0.18	0.01	0.29	0.02	22	15	31	$\sim$	$\sim$	0.67	0.13
8×10^5	8	6	4	$\sim$	$\sim$	$\sim$	$\sim$	8	1	15	$\sim$	$\sim$	$\sim$	$\sim$

T2 and **T3** trajectories in a common frame

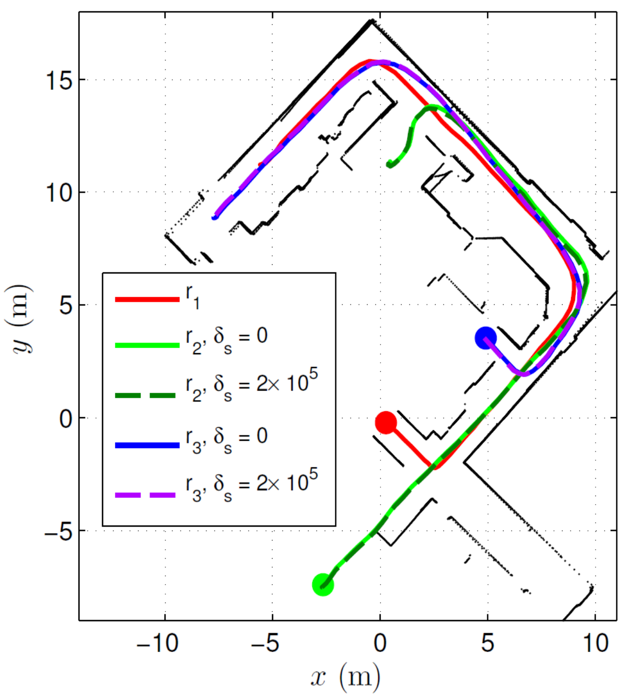

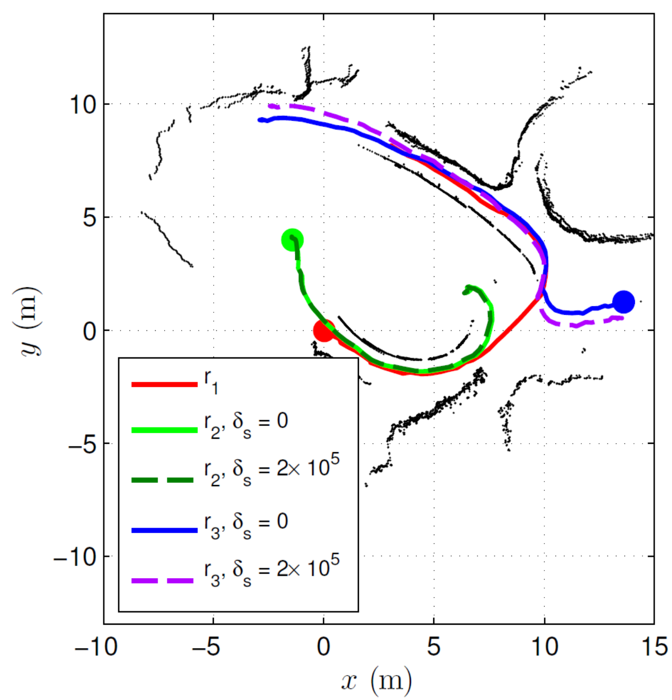

Results: Saliency Thresholding

T2 and **T3** robots mapping in a computed common frame

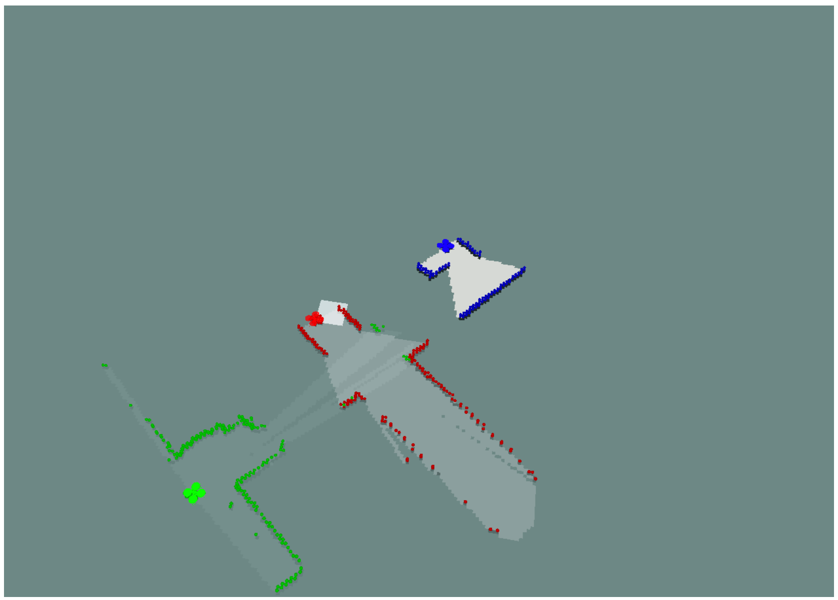

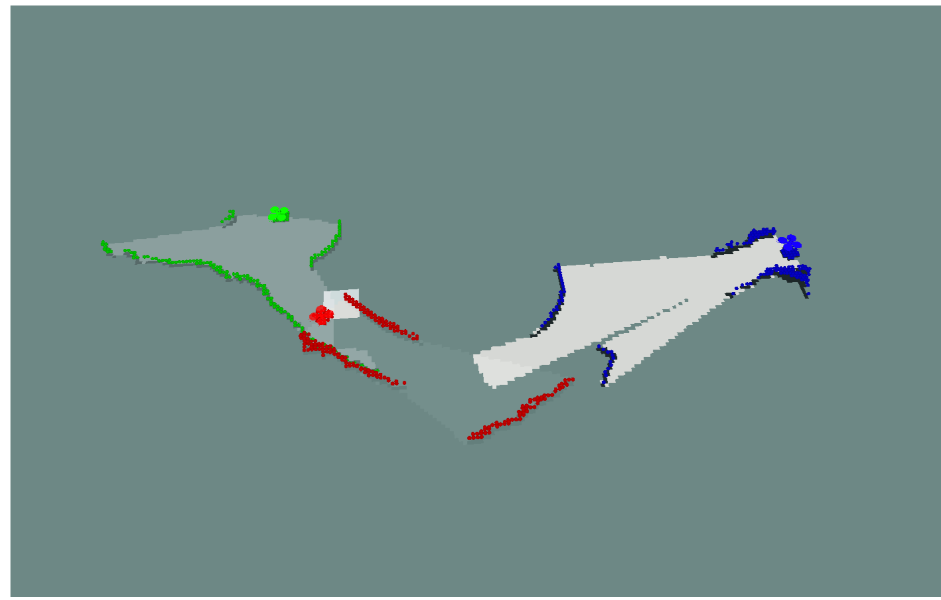

Results: Sharing Frequency and Run Time

Capacity constrained networking

- ~34 kB per scan
- 4 Hz sharing limit with n=3
- 1 Hz sharing limit with n=6

[Jun, et al., IEEE Wireless Communications 2003]

Robot sharing frequencies

		$\delta_s = 0$			$\delta_s = 2 \times 10^5$		
Robot	Duration (s)	Shared Scans	Max (Hz)	Mean (Hz)	Shared Scans	Max (Hz)	Mean (Hz)
T2: r_1	37.4	75	2.08	2.00	22	1.01	0.59
T2: r_2	39.0	77	2.02	1.97	26	1.28	0.67
T2: r_3	32.5	65	2.02	2.00	23	0.95	0.71
T3: r_1	35.5	71	2.06	2.00	26	0.98	0.73
T3: r_2	27.6	55	1.99	1.99	18	1.31	0.65
T3: r_3	37.4	74	1.98	1.98	36	1.20	0.96

- Without thresholding saliency, network capacity is not reached
- Thresholding causes a reduction in both mean and max sharing frequencies

Percentage of total run time devoted to individual algorithmic steps

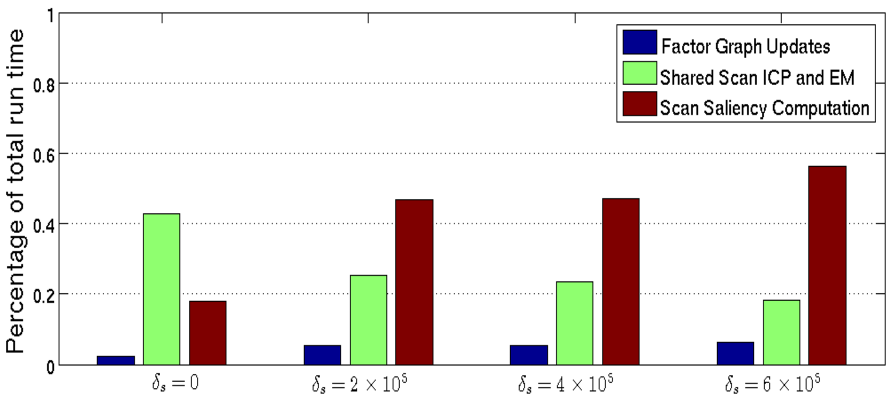

Trial **T3**

- Scan saliency computation requires the same amount of time regardless of the number of shared observations
- Therefore run time was decreased by 46.4% by discarding the bottom 60.0% of salient scans

Conclusions

Experimental analysis of multi-robot data association framework

- Laser scan autocovariance as a measure of saliency
- Subsampling by saliency reduces complexity, mildly diminishes transformation accuracy
- With three robots, implementation is not constrained by network capacities