Inferring and Enforcing Relative Constraints in SLAM

Kris Beevers and Wes Huang

Department of Computer Science
Rensselaer Polytechnic Institute
{beevek,whuang}@cs.rpi.edu

July 16, 2006

Overview

- **Goal:** exploit prior knowledge of environment to improve SLAM
- Example: many indoor environments are "mostly" rectilinear
  - *Linear equality constraints* on (r, θ) lines, e.g.: $\theta_1 = \theta_2 + \frac{\pi}{2}$

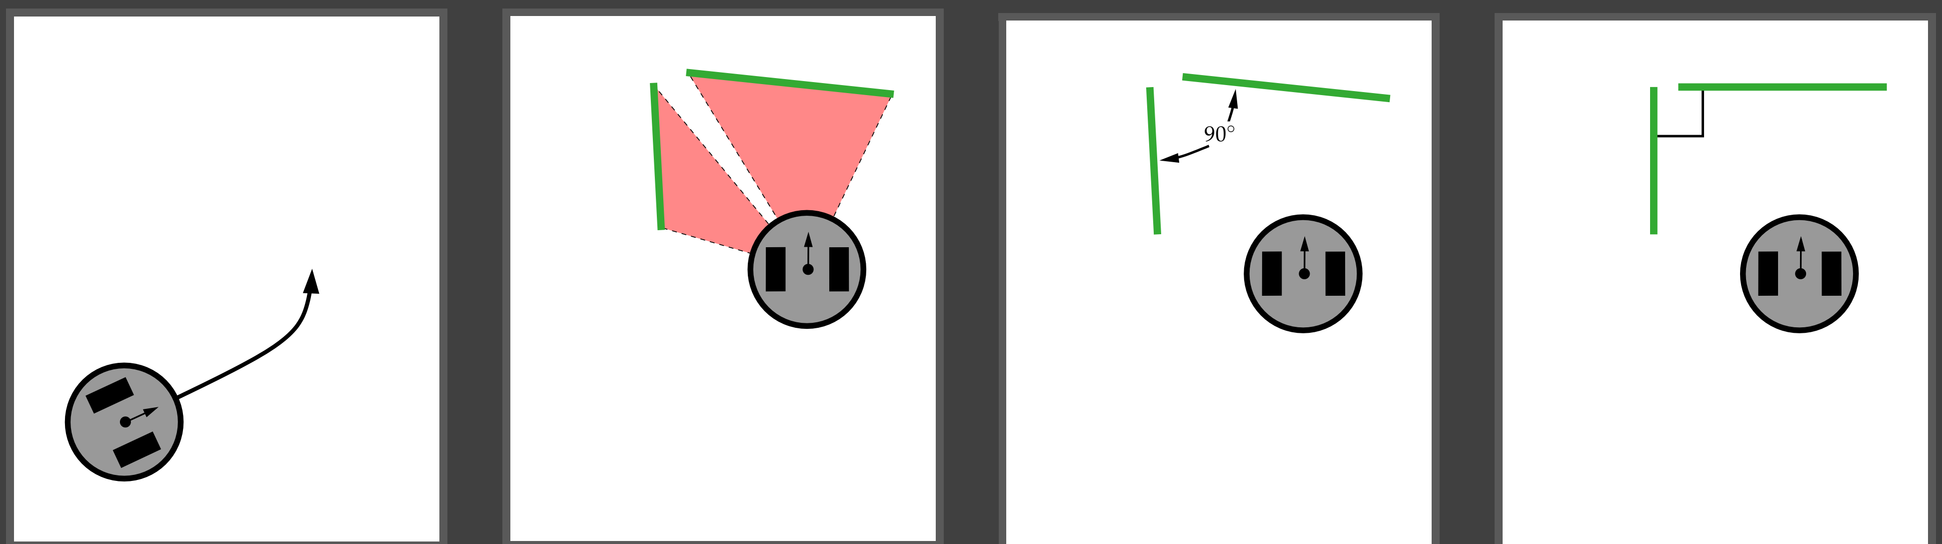

predict | sense/update | infer constraints | enforce constraints

- **Inference**: *when* to apply a relative constraint to the map
- **Enforcement**: *how* to apply constraints in (RBPF) SLAM

Previous work

- Constraint inference:
  - Loop closing (Newman, 1999)
  - Gating approaches (Rodriguez-Losada et al., 2006)
- Enforcing constraints on the map in EKF-SLAM:
  - Treat constraints as zero-uncertainty measurements (Durrant-Whyte, 1988; Wen and Durrant-Whyte, 1992)
  - Project unconstrained estimate onto constraint surface (Simon and Chia, 2002; Simon and Simon, 2003)
- Relative maps: apply constraints on relationships to ensure consistency (Csorba and Durrant-Whyte, 1997; Newman, 1999; Deans and Hebert, 2000)

Rao-Blackwellized particle filtering SLAM

- Estimate map $M = \{x_1, \ldots, x_n\}$ and robot trajectory s^t from measurements z^t, control inputs u^t, and correspondences n^t

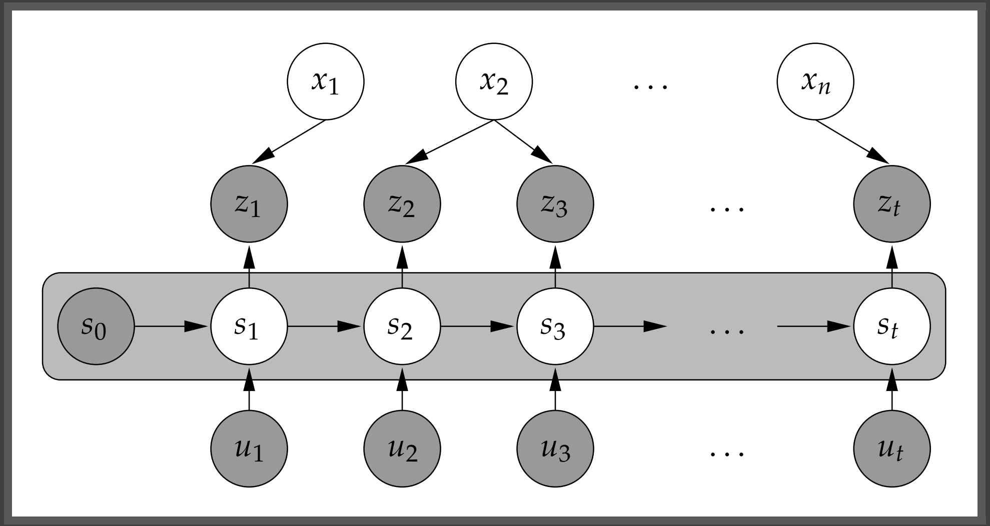

- Landmarks independent conditioned on trajectory
- Robot's trajectory s^t "d-separates" the landmark variables $\{x_i\}$

$$\underbrace{p(s^t, M | z^t, u^t, n^t)}_{\text{posterior}} = \underbrace{p(s^t | n^t, z^t, u^t)}_{\text{posterior over trajectories}} \prod_{i=1}^{n} \underbrace{p(x_i | s^t, n^t, z^t)}_{\text{posterior over landmark } i}$$

- Estimate $p(s^t)$ by N samples, $p(x_i | s^t)$ by n small EKFs: $O(Nn)$

Structured environments

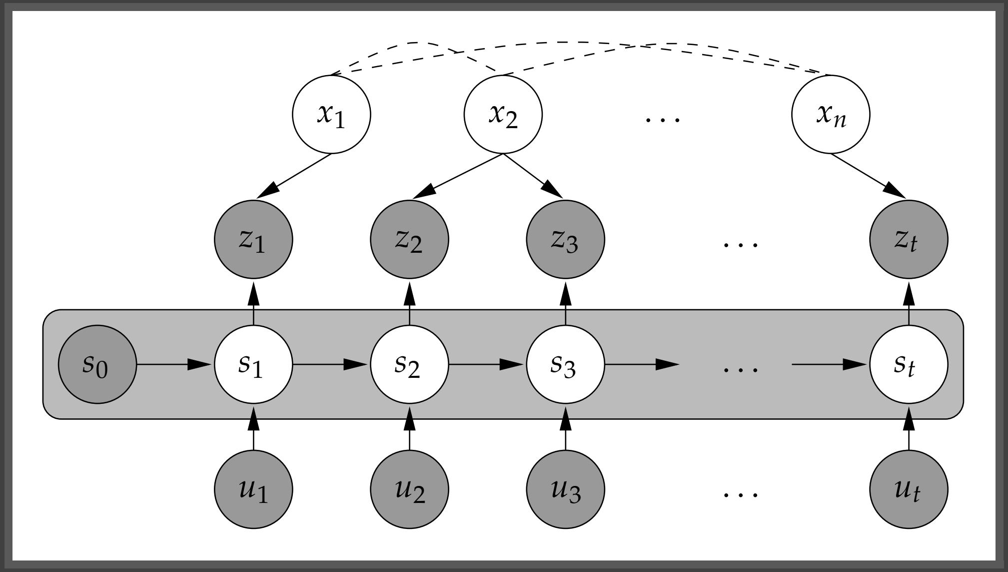

- RBPF SLAM **assumption**: environment is unstructured
  - Landmarks are *randomly and independently distributed*
- **But**: many environments exhibit structure (e.g., indoors)
  - Architects do not (usually) throw darts
- Correlation between landmarks arises because of structure
  - This breaks the RBPF factorization!

Extending the model

- **Approach:** model the correlations
- **Model**: relationship between x_i and x_j is parameterized by $c_{i,j}$
- Assume environment structure takes on one of a few forms
  - Space of (structured) landmark relationships is small and discrete
  - Rectilinearity: $c_{i,j} \in \{0, 90, 180, 270, \star\}$
- Do inference in parameter space
- Treat the relationships as constraints to be enforced

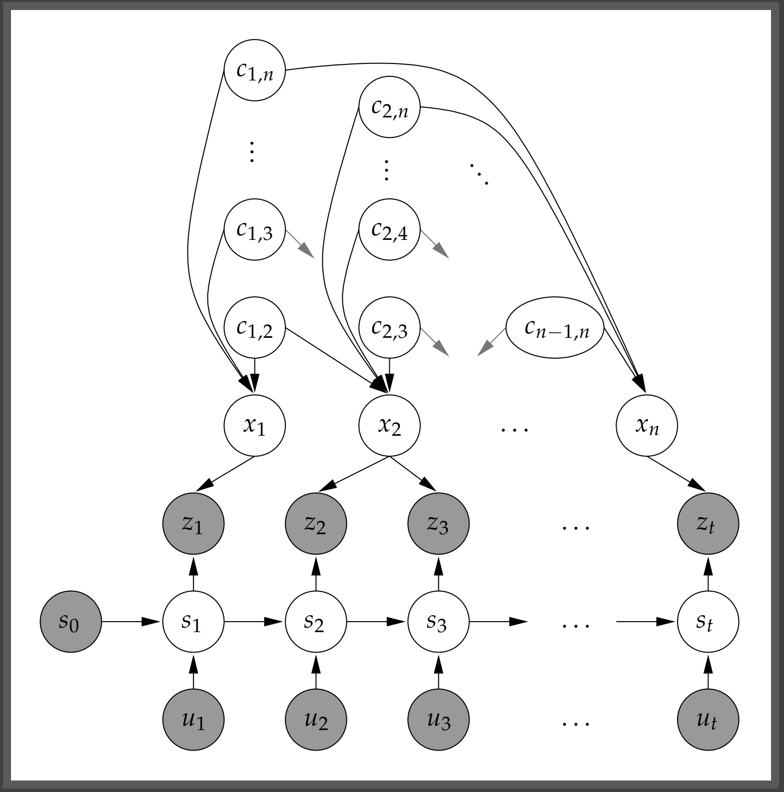

Superlandmark filter

- Given $\{c_{i,j}\}$, how to enforce relative constraints on the map?

- Idea: group sets of constrained landmarks into "superlandmarks"

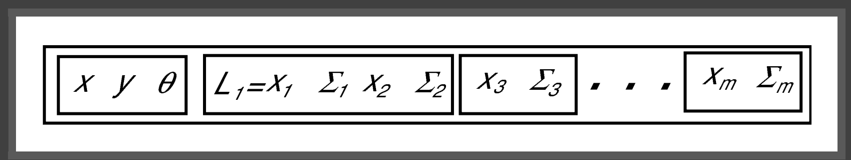

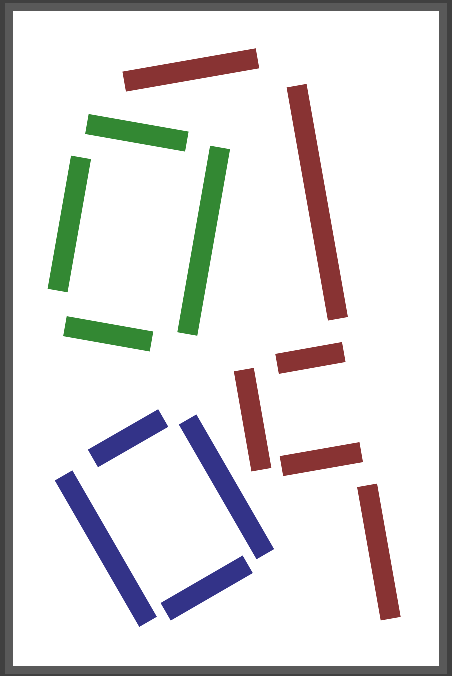

- Estimate each superlandmark in a particle's map with an EKF
- Superlandmarks are independent conditioned on s^t

- Apply EKF-based constraint enforcement techniques to each superlandmark (Durrant-Whyte, 1988; Simon and Chia, 2002)

- **Not a good idea by itself!** $O(Nn^3)$

Reduced state formulation

- A simple improvement for linear equality constraints:
  - Instead of: $L_i = \{r_1, \theta_1, r_2, \theta_2, \ldots\}$
  - Use: $L_i = \{\theta_1, r_1, r_2, \ldots\}$ since $\theta_i = g_i(c_{1,i}; \theta_1)$
- Superlandmark filtering over this is still $O(Nn^3)$
- Instead: apply Rao-Blackwellization to the reduced state

Rao-Blackwellized constraint filter

- Divide map into constrained and unconstrained variables: $M = \{M^c, M^u\}$
  - Rectilinearity: for a superlandmark $L_i = \{\theta, r_1, r_2, \ldots\}$: $M^c = \{\theta\}$ and $M^u = \{r_1, r_2, \ldots\}$
- Sample the trajectory *and the constrained variables*, estimate unconstrained variables with EKFs

$$\boxed{\underbrace{p(s^t, M | z^t, u^t, n^t)}_{\text{posterior}} = \underbrace{p(s^t, M^c | n^t, z^t, u^t)}_{\text{trajectories/constrained}} \prod_i \underbrace{p(x_i^u | s^t, M^c, n^t, z^t)}_{\text{unconstrained vars. of LM } i}}$$

- Computational complexity is now $O(Nn)$ (same as normal RBPF)
- Tricky details:
  - How to sample constrained variables?
  - Adding new constraints between previously-constrained landmarks

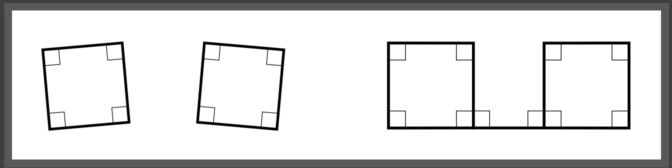

Inference of constraints

- Should landmarks x_i, x_j be constrained with respect to each other?

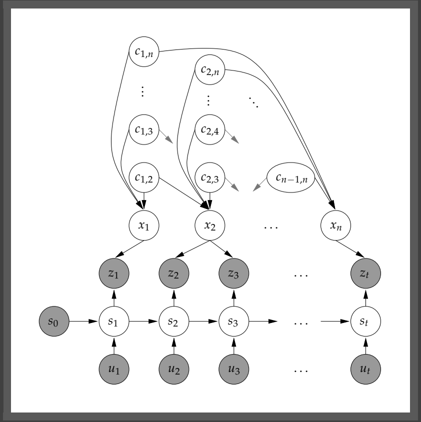

- An inference problem on constraint parameters $c_{i,j}$

- Gating (Rodriguez-Losada et al., 2006):
  - Constrain if $|\theta_i - \theta_j| \approx \{0°, 90°, 180°, 270°\}$
  - Ignores confidence in landmark estimates

- Our approach:
  - Compute PMF over constraint parameter space (e.g., $\{0, 90, 180, 270, \star\}$) at landmark initialization time
  - Sample from the PMF for each RBPF particle
  - Particles with incorrectly constrained landmarks will eventually be resampled

Results

- Implementation of rectilinearity constraints on top of sparse-sensing SLAM (Beevers and Huang, 2006)

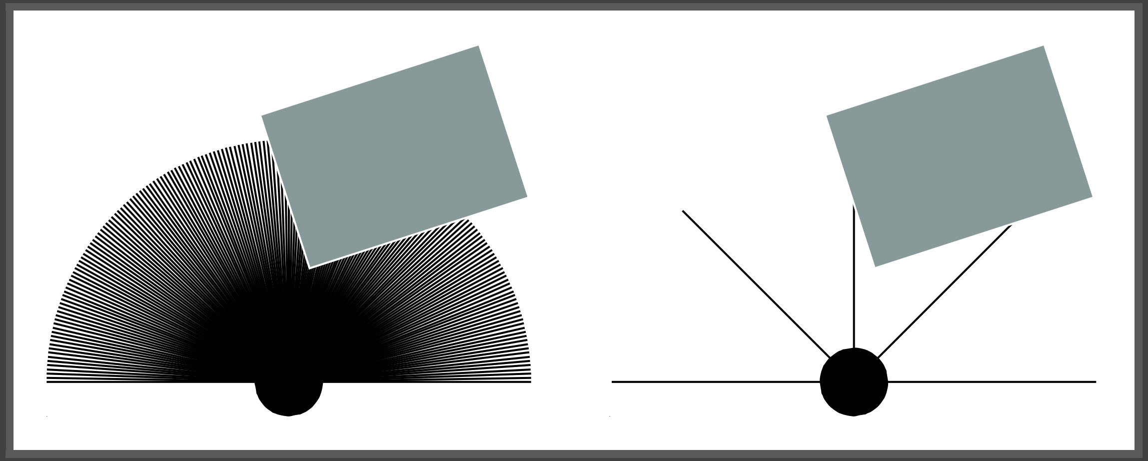

  - Lack of data requires delayed landmark initialization
  - High uncertainty, many particles
- Incorporating rectilinearity constraints when applicable:
  - Improves filter consistency
  - Allows mapping with many fewer particles

Real-world results: USC SAL

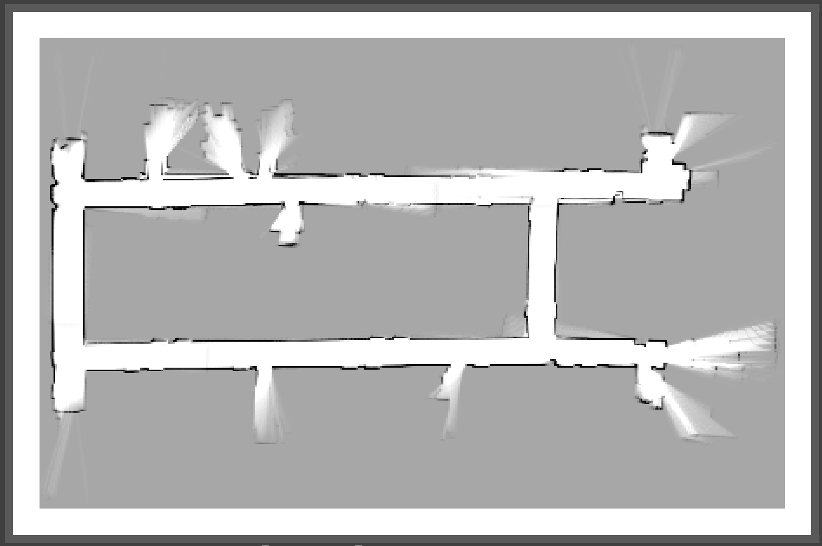

unconstrained
100 particles
32.02 sec

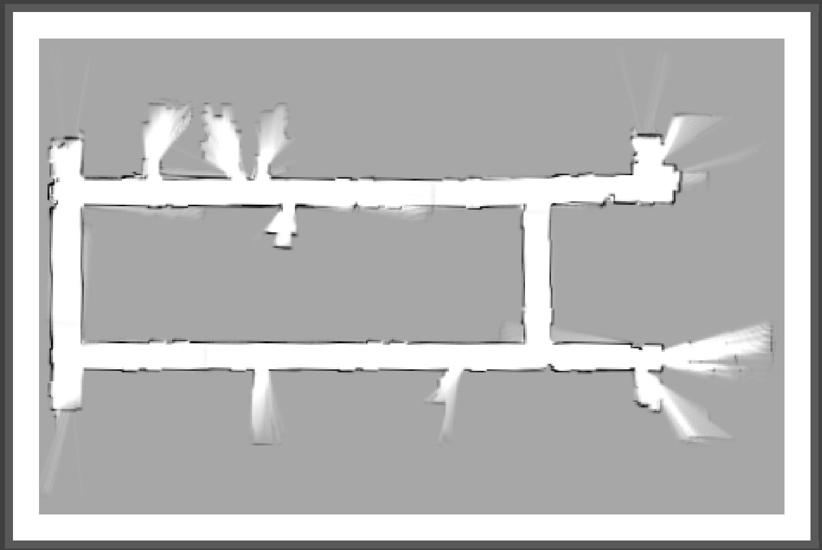

constrained
20 particles
11.24 sec

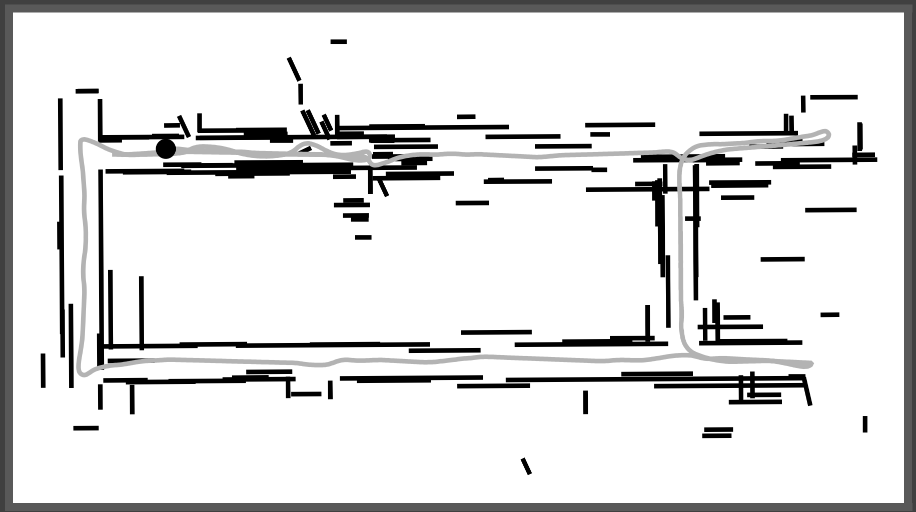

Data from Radish courtesy Andrew Howard

Real-world results: CMU NSH

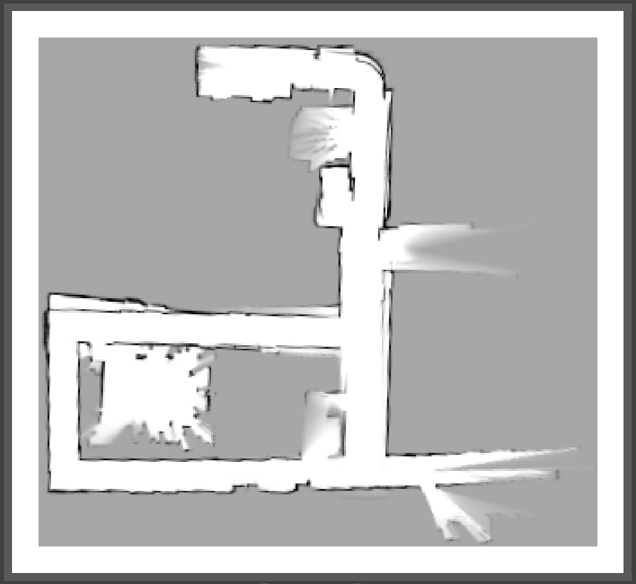

unconstrained
600 particles
268.44 sec

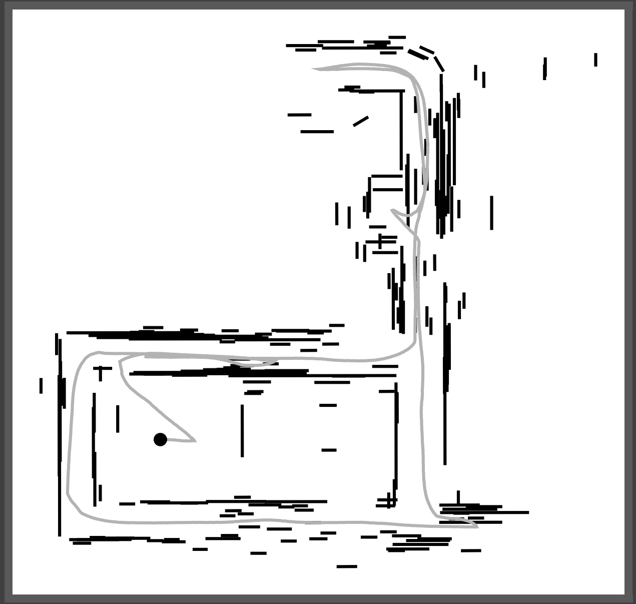

Data from Radish courtesy Nick Roy

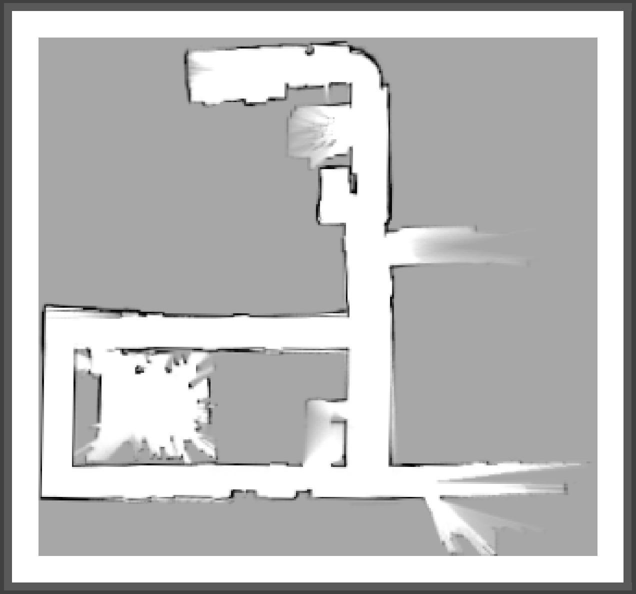

constrained
40 particles
34.77 sec

Discussion

- Caveat: conditioning on constrained variable values is sensitive to covariance estimation inaccuracies
  - Inaccurate covariance results in landmark drift
  - Probably not an issue with more data (e.g., laser rangefinder)
- More variables being estimated by particles — but we need fewer particles?
  - Constraints reduce the DOF of the map
  - Reduced state: only one sampled variable per superlandmark

Simulation results: consistency

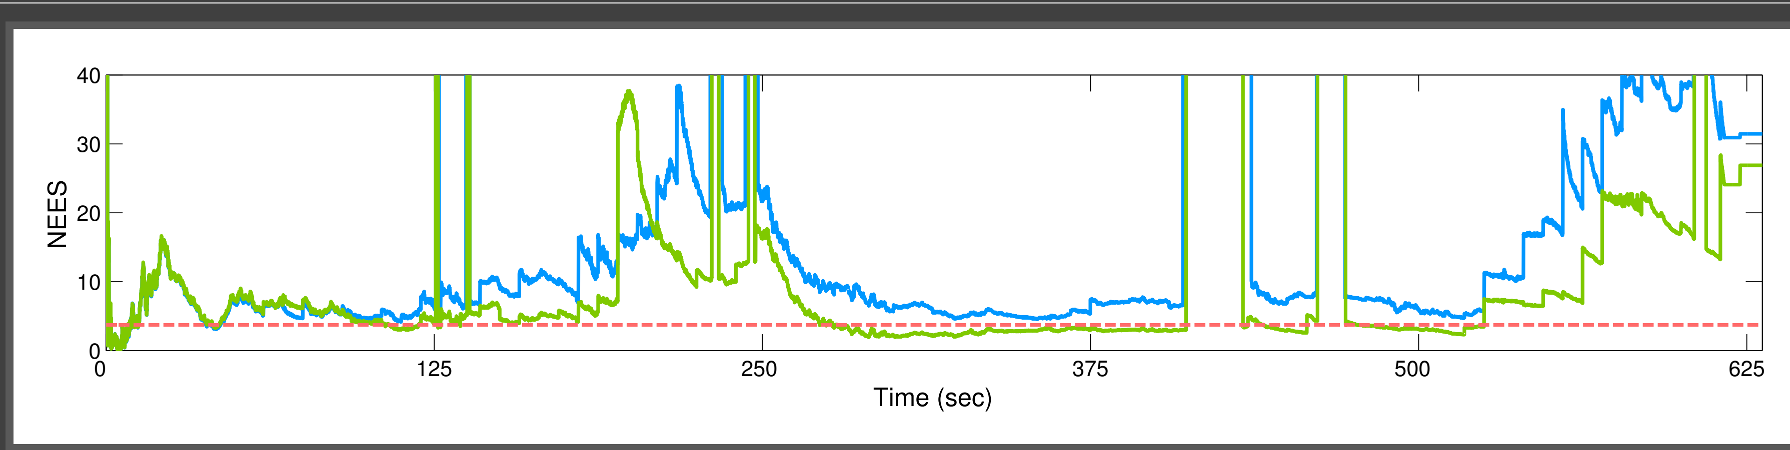

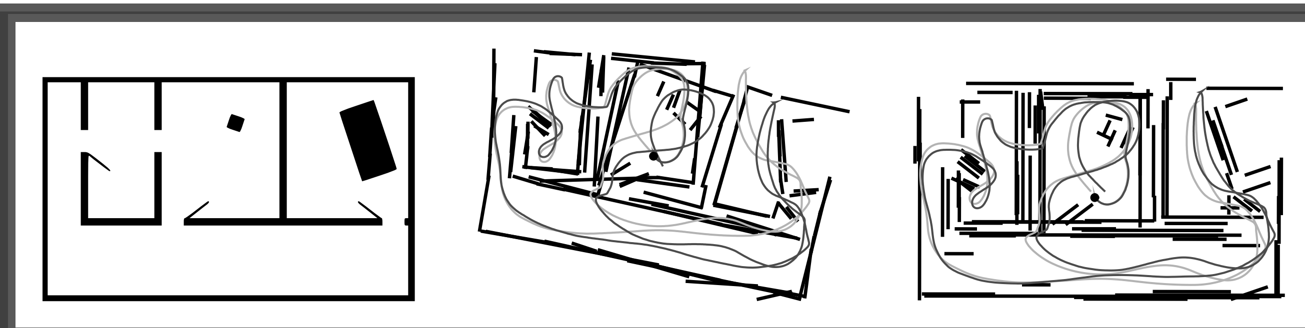

- Inconsistency: filter significantly underestimates its own error
- Bailey et al. (2006): normalized estimation error squared (NEES) as a measure of RBPF SLAM consistency: $(s_t - \hat{s}_t)\hat{\Sigma}_{s_t}^{-1}(s_t - \hat{s}_t)^T$
- 200 particles, 50 Monte Carlo trials

Summary

- RBPF SLAM algorithm that models structural relationships of landmarks as parameterized linear equality constraints
  - Perform inference in the space of constraints
  - Mechanism for enforcing constraints in RBPF
- Results:
  - Improved filter consistency
  - Successful mapping (real and simulated) with many fewer particles than unconstrained SLAM
- Future work:
  - Eliminating landmark drift due to poor covariance estimation
  - Different types of constraints (e.g., inequality — see paper)

Thank you!

Some details

- Sampling constrained variables ("particlization"):
  - Combine unconstrained estimates from all landmarks in the superlandmark
  - Example: $\hat{\theta}_1 = -1^\circ, \hat{\theta}_2 = 90^\circ$ with constraint $\theta_2 = \theta_1 + 90^\circ$
  - Compute maximum likelihood mean $\hat{\theta}$ and variance $\hat{\sigma}_\theta$
  - Sample from $\mathcal{N}(\hat{\theta}, \hat{\sigma}_\theta)$
- Adding constraints between superlandmarks ("reconditioning"):

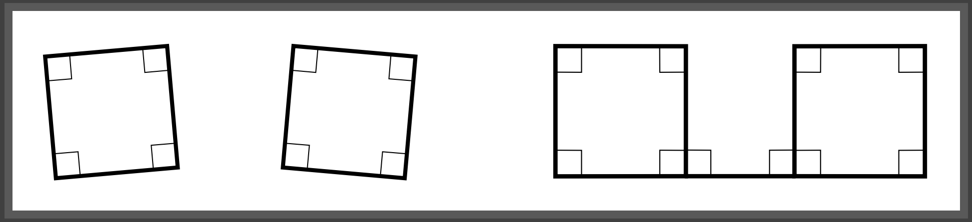

  - Keep measurements since particlization in an "accumulator"
  - "Rewind" to time of particlization, condition on new values of constrained variables, apply accumulated measurements

References

T. Bailey, J. Nieto, and E. Nebot. Consistency of the FastSLAM algorithm. In *IEEE Intl. Conf. on Robotics and Automation*, pages 424–427, 2006.

Kristopher R. Beevers and Wesley H. Huang. SLAM with sparse sensing. In *IEEE Intl. Conf. on Robotics and Automation*, pages 2285–2290, 2006.

M. Csorba and H.F. Durrant-Whyte. New approach to map building using relative position estimates. *SPIE Navigation and Control Technologies for Unmanned Systems II*, 3087(1):115–125, 1997.

M.C. Deans and M. Hebert. Invariant filtering for simultaneous localization and mapping. In *IEEE Intl. Conf. on Robotics and Automation*, pages 1042–1047, 2000.

H.F. Durrant-Whyte. Uncertain geometry in robotics. *IEEE Journal of Robotics and Automation*, 4(1):23–31, 1988.

P.M. Newman. *On the structure and solution of the simultaneous localization and mapping problem*. PhD thesis, University of Sydney, Australia, 1999.

D. Rodriguez-Losada, F. Matia, A. Jimenez, and R. Galan. Consistency improvement for SLAM – EKF for indoor environments. In *IEEE Intl. Conf. on Robotics and Automation*, pages 418–423, 2006.

D. Simon and T.L. Chia. Kalman filtering with state equality constraints. *IEEE Transactions on Aerospace and Electronic Systems*, 39:128–136, 2002.

D. Simon and D.L. Simon. Aircraft turbofan engine health estimation using constrained Kalman filtering. In *ASME Turbo Expo*, 2003.

W. Wen and H.F. Durrant-Whyte. Model-based multi-sensor data fusion. In *IEEE Intl. Conf. on Robotics and Automation*, pages 1720–1726, 1992.