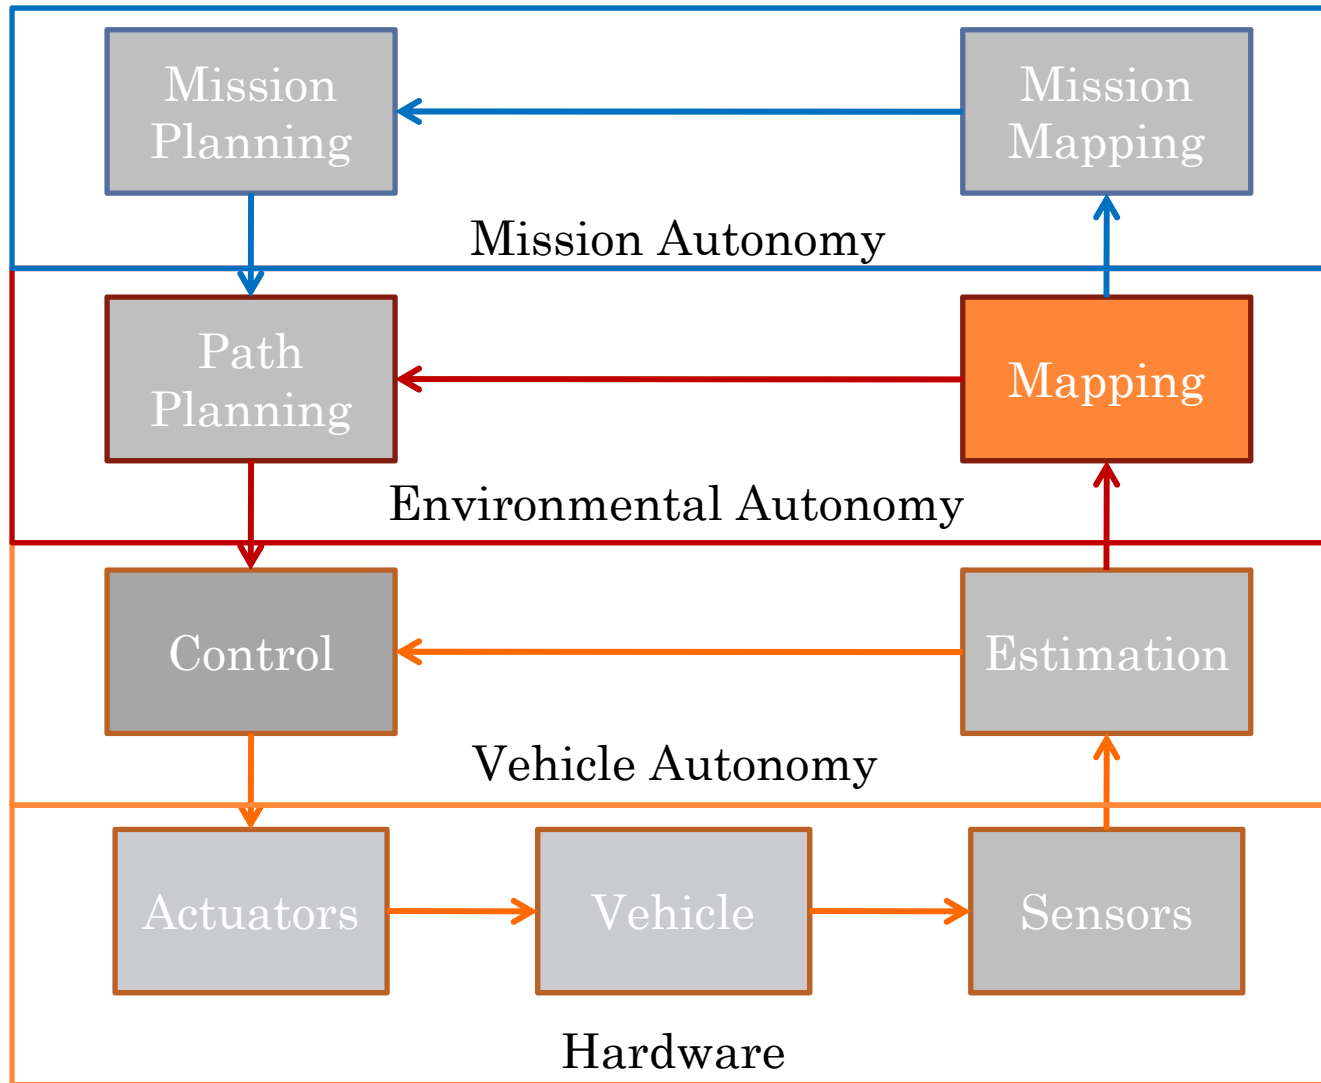


ME 597: AUTONOMOUS MOBILE ROBOTICS SECTION 8 – MAPPING III

Prof. Steven Waslander

COMPONENTS



OUTLINE

- The GraphSLAM algorithm
 - Derivation of feature based optimization problem
 - Derivation of scan based optimization problem
 - Discussion of solution methods
 - Implementation and Results

TWO MAIN SLAM APPROACHES

○ Online SLAM

- Filter version of the SLAM problem, maximize

$$p(x_t | y_{1:t}, u_{1:t})$$

- Process new information as it is received
- Generate current best estimate, rely on Markov assumption and linearity to trust that this is the best you can do, and use the solution in subsequent steps
- EKF SLAM, FastSLAM Occupancy Grid SLAM, etc.

○ Full SLAM

- Smoothing version of the SLAM problem, maximize

$$p(x_{0:t} | y_{1:t}, u_{1:t})$$

- Store all information as collected, only resolve into poses and map when needed
- Work on all information, allows for re-linearization during the optimization process
- Can resolve correspondence as well, allowing for a more robust solution

FULL SLAM PROBLEM

○ Full SLAM - Features

- Simultaneously determine the robot pose history and static feature locations in the environment.

$$x_t^r = \begin{pmatrix} X_t \\ Y_t \\ Z_t \\ \phi_t \\ \theta_t \\ \psi_t \end{pmatrix}, \quad m_i = \begin{pmatrix} m_{i,X} \\ m_{i,Y} \\ m_{i,Z} \end{pmatrix}, \quad m = \begin{pmatrix} m_1 \\ \vdots \\ m_M \end{pmatrix}, \quad x_{0:t} = \begin{pmatrix} x_0^r \\ x_1^r \\ \vdots \\ x_t^r \\ m \end{pmatrix}, \quad x_t = \begin{pmatrix} x_t^r \\ m \end{pmatrix}$$

Robot State
at time t

i^{th}
feature

Feature
map

Full
state

Full state
at time t

FULL SLAM PROBLEM

- Available information

- Inputs and Motion model


$$u_{0:t}, \quad x_t^r = g(x_{t-1}^r, u_t) + \delta_t$$

- Measurements and measurement model

$$y_{1:t}, \quad y_t = h(x_t^r, m) + \varepsilon_t$$

- Again, we'll assume good correspondence information, but this is an important part of GraphSLAM, can include correspondence as part of the optimization

ILLUSTRATION OF THE CONSTRAINT GRAPH



GRAPHSLAM OPTIMIZATION

- We are interested in finding the maximum likelihood state estimate

$$\max_{x_{0:t}} p(x_{0:t} | y_{1:t}, u_{1:t})$$

- Apply Bayes rule to separate out the current measurements

$$\begin{aligned} p(x_{0:t} | y_{1:t}, u_{1:t}) &= \eta p(y_t | x_{0:t}, u_{1:t}) p(x_{0:t} | y_{1:t-1}, u_{1:t}) \\ &= \eta p(y_t | x_t) p(x_{0:t} | y_{1:t-1}, u_{1:t}) \end{aligned}$$

GRAPHSLAM OPTIMIZATION

- Next, separate out the motion through factoring of the probabilities of second term, since y_t is not present

$$\begin{aligned} p(x_{0:t} | y_{1:t-1}, u_{1:t}) \\ &= p(x_t^r | x_{0:t-1}, u_{1:t}) p(x_{0:t-1} | y_{1:t-1}, u_{1:t}) \\ &= p(x_t^r | x_{t-1}, u_t) p(x_{0:t-1} | y_{1:t-1}, u_{1:t}) \end{aligned}$$

- These steps we repeat until the beginning of time to get

$$\begin{aligned} p(x_{0:t} | y_{1:t}, u_{1:t}) &= \eta p(x_0) \prod_{\tau=1}^t p(x_\tau^r | x_{\tau-1}^r, u_\tau) p(y_\tau | x_\tau) \\ &= \eta p(x_0) \prod_{\tau=1}^t \left(p(x_\tau^r | x_{\tau-1}^r, u_\tau) \prod_i p(y_\tau^i | x_\tau) \right) \end{aligned}$$

- If there is no prior information about the map, use $p(x^r_0)$

GRAPHSLAM OPTIMIZATION

- We can redefine our optimization problem as

$$\max_{x_{0:t}} p(x_{0:t} | y_{1:t}, u_{1:t})$$



$$\max_{x_{0:t}} \eta p(x_0) \prod_{\tau=1} \left(p(x_{\tau}^r | x_{\tau-1}^r, u_{\tau}) \prod_i p(y_{\tau}^i | x_{\tau}) \right)$$



$$\min_{x_{0:t}} -\ln \left(\eta p(x_0) \prod_{\tau=1} \left(p(x_{\tau}^r | x_{\tau-1}^r, u_{\tau}) \prod_i p(y_{\tau}^i | x_{\tau}) \right) \right)$$

GRAPHSLAM OPTIMIZATION

- We can redefine our optimization problem as

$$\min_{x_{0:t}} -\ln \left(\eta p(x_0) \prod_{\tau=1} \left(p(x_{\tau}^r | x_{\tau-1}^r, u_{\tau}) \prod_i p(y_{\tau}^i | x_{\tau}) \right) \right)$$



$$\min_{x_{0:t}} J = \text{const.} - \ln(p(x_0))$$

$$- \sum_{\tau=1}^t \left(\ln(p(x_{\tau}^r | x_{\tau-1}^r, u_{\tau})) \right) - \sum_{\tau=1}^t \sum_i \ln(p(y_{\tau}^i | x_{\tau}))$$

GRAPHSLAM OPTIMIZATION

- The assumption about additive Gaussian noise and disturbances means that the motion and measurement models can be expressed as Gaussian distributions

- Motion

$$p(x_t^r | x_{t-1}^r, u_t) = \eta e^{-\frac{1}{2} [x_t^r - g(x_{t-1}^r, u_t)]^T R^{-1} [x_t^r - g(x_{t-1}^r, u_t)]}$$

- Measurement

$$p(y_t^i | x_t) = \eta e^{-\frac{1}{2} [y_t^i - h(x_t)]^T Q^{-1} [y_t^i - h(x_t)]}$$

- Prior $p(x_0) = \eta e^{-\frac{1}{2} [x_0 - \mu_0]^T \Sigma_0^{-1} [x_0 - \mu_0]}$

$$\mu_0 = 0, \quad \Sigma_0 = 0, \quad \Sigma_0^{-1} = \infty I$$

GRAPHSLAM OPTIMIZATION

- The negative log likelihoods therefore all take the *Mahalonobis distance* form

- Motion

$$-\ln p(x_t^r | x_{t-1}^r, u_t) = \text{const.} + [x_t^r - g(x_{t-1}^r, u_t)]^T R^{-1} [x_t^r - g(x_{t-1}^r, u_t)]$$

- Measurement

$$-\ln p(y_t^i | x_t) = \text{const.} + [y_t^i - h(x_t)]^T Q^{-1} [y_t^i - h(x_t)]$$

- Prior

$$-\ln p(x_0) = \text{const.} + [x_0 - \mu_0]^T \Sigma_0^{-1} [x_0 - \mu_0]$$

GRAPHSLAM OPTIMIZATION

- The final form of the optimization is now

$$\begin{aligned} \min_{z_{0:t}} J = & \text{const.} + \left[x_0 - \mu_0 \right]^T \Sigma_0^{-1} \left[x_0 - \mu_0 \right] \\ & + \sum_{\tau=1}^t \left[x_{\tau}^r - g(x_{\tau-1}^r, u_{\tau}) \right]^T R^{-1} \left[x_{\tau}^r - g(x_{\tau-1}^r, u_{\tau}) \right] \\ & + \sum_{\tau=1}^t \sum_i \left[y_{\tau}^i - h(x_{\tau},) \right]^T Q^{-1} \left[y_{\tau}^i - h(x_{\tau},) \right] \end{aligned}$$

- This is an unconstrained nonlinear optimization problem, which now needs to be solved somehow.
 - There is a lot of structure to the problem, because of the sequential nature of the motion constraints and the measurement of features at only a few instances in time.

GRAPH CONSTRAINTS

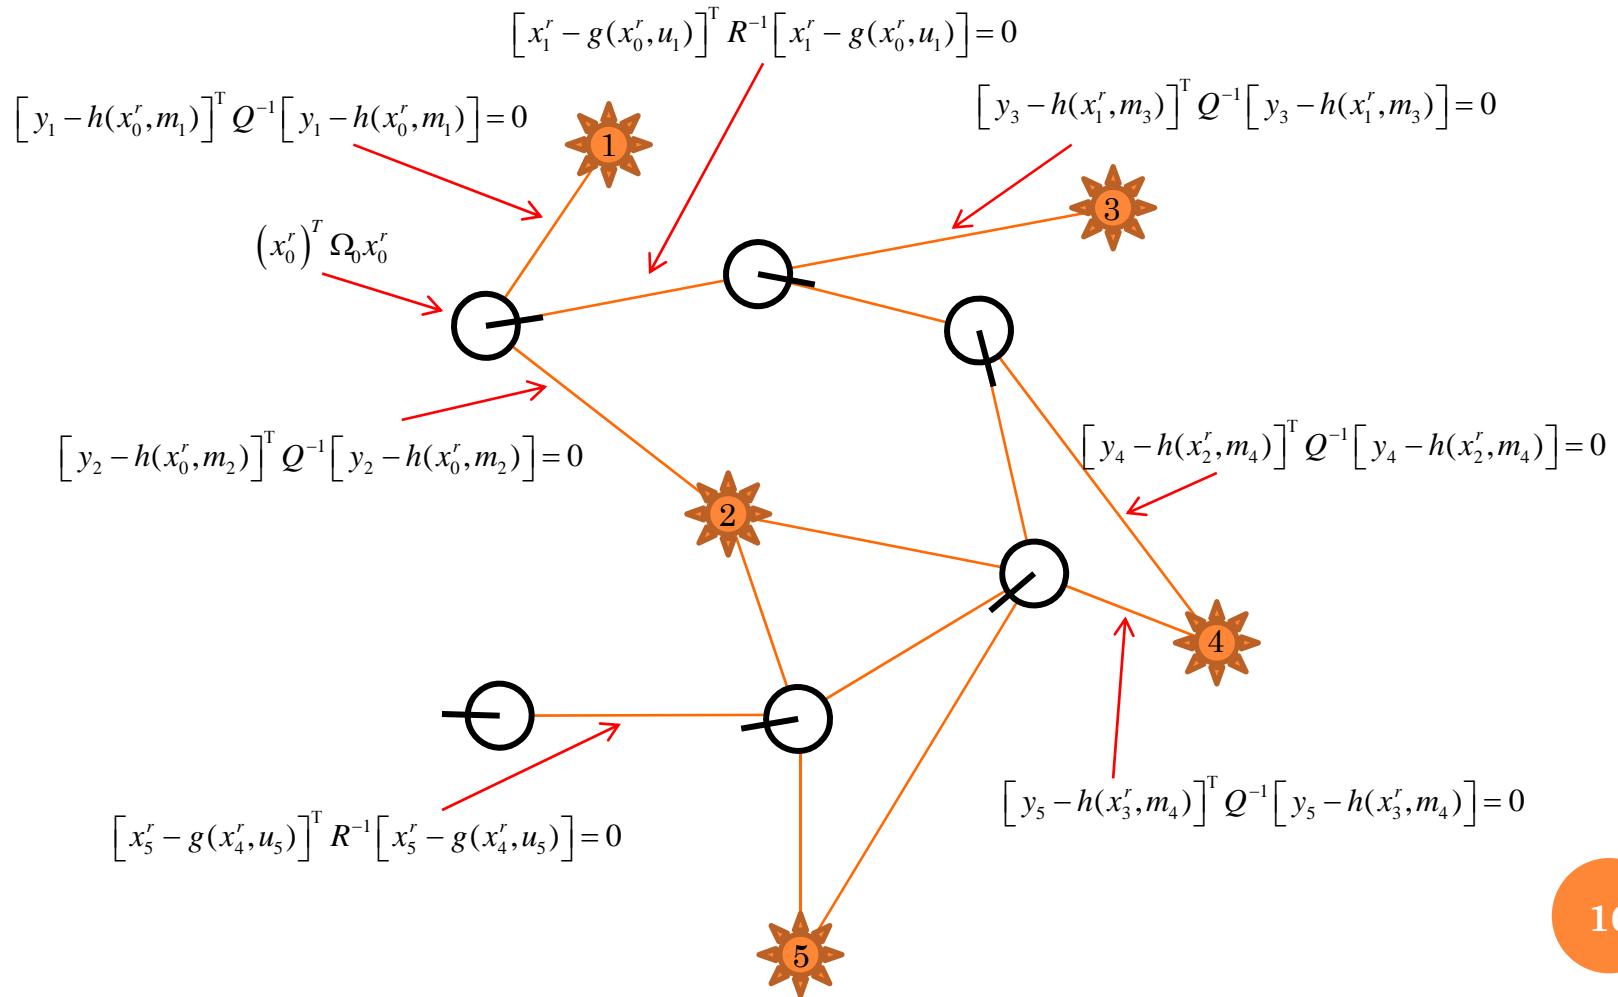
- The constraints on the graph can now be thought of in a least squares sense.
 - Over-determined set of constraints, optimization aims to minimize the total violation of the full set of constraints
 - Can be considered a weighted distance minimization
 - Errors minimized together based on inverse of covariance weighting (information matrix)
 - Motion

$$\left[x_t^r - g(x_{t-1}^r, u_t) \right]^T R^{-1} \left[x_t^r - g(x_{t-1}^r, u_t) \right] = 0$$

- Measurement

$$\left[y_t - h(x_t, c_t) \right]^T Q^{-1} \left[y_t - h(x_t, c_t) \right] = 0$$

ILLUSTRATION OF THE CONSTRAINT GRAPH



GRAPHSLAM OPTIMIZATION

- For standard nonlinear optimization packages, you must provide

- Cost function

$$J = \text{const.} + [x_0 - \mu_0]^T \Sigma_0^{-1} [x_0 - \mu_0] + \sum_{\tau=1}^t [x_\tau^r - g(x_{\tau-1}^r, u_\tau)]^T R^{-1} [x_\tau^r - g(x_{\tau-1}^r, u_\tau)] \\ + \sum_{\tau=1}^t \sum_i [y_\tau^i - h(x_\tau)]^T Q^{-1} [y_\tau^i - h(x_\tau)]$$

- Gradient function

$$\frac{\partial J}{\partial x_0} = [x_0 - \mu_0]^T \Sigma_0^{-1} + [x_1^r - g(x_0^r, u_1)]^T R^{-1} \frac{\partial}{\partial x_0} [g(x_0^r, u_1)]$$

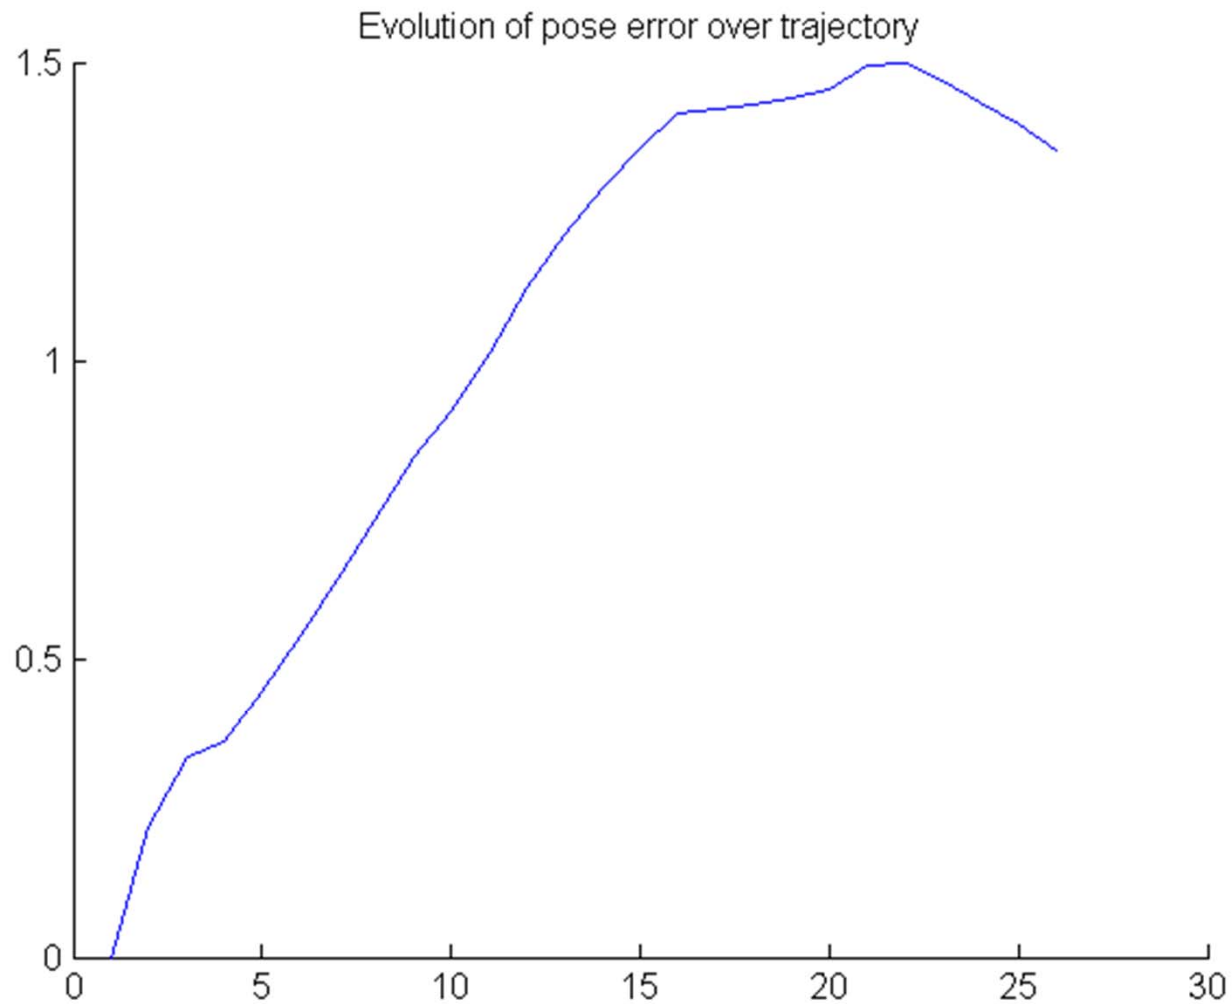
$$\frac{\partial J}{\partial x_t^r} = -[x_{t+1}^r - g(x_t^r, u_{t+1})]^T R^{-1} \frac{\partial}{\partial x_t^r} [g(x_t^r, u_{t+1})] + [x_t^r - g(x_{t-1}^r, u_t)]^T R^{-1} \\ - \sum_i [y_t^i - h(x_t)]^T Q^{-1} \frac{\partial}{\partial x_t^r} [h(x_t)]$$

$$\frac{\partial J}{\partial x_t^m} = - \sum_i [y_t^i - h(x_t)]^T Q^{-1} \frac{\partial}{\partial x_t^m} [h(x_t)]$$

- Initial Estimate of complete state (from odometry, other sensors)

$$\tilde{x}_{0:t}$$

GRAPHSLAM PRELIMINARY RESULTS



GRAPHSLAM

- GraphSLAM by Thrun and Montemerlo [2006]
 - Many interesting customizations to make optimization tractable
 - Linearization of models to form locally quadratic problem
 - Factorization of map into robot poses to reduce graph size
 - Scan points used as features with correspondence updated inside optimization
 - Full details in Chap 11 of Probabilistic Robotics
 - Summarized in extra slides at the end of this presentation
 - Warning: slightly different notation used

OUTLINE

- The GraphSLAM algorithm
 - Derivation of feature based optimization problem
 - Derivation of scan based optimization problem
 - Discussion of solution methods
 - Implementation and Results

GRAPHSLAM WITH SCAN REGISTRATION

○ GraphSLAM – Scan Registration

- No map elements are included in the state vector.
- Instead, all scans are converted into relative pose measurements through registration


$$x_t^r = \begin{pmatrix} X_t \\ Y_t \\ Z_t \\ \phi_t \\ \theta_t \\ \psi_t \end{pmatrix}, \quad x_{0:t}^r = \begin{pmatrix} x_0^r \\ x_1^r \\ \vdots \\ \vdots \\ x_t^r \end{pmatrix}$$

Robot State
at time t

Full
state

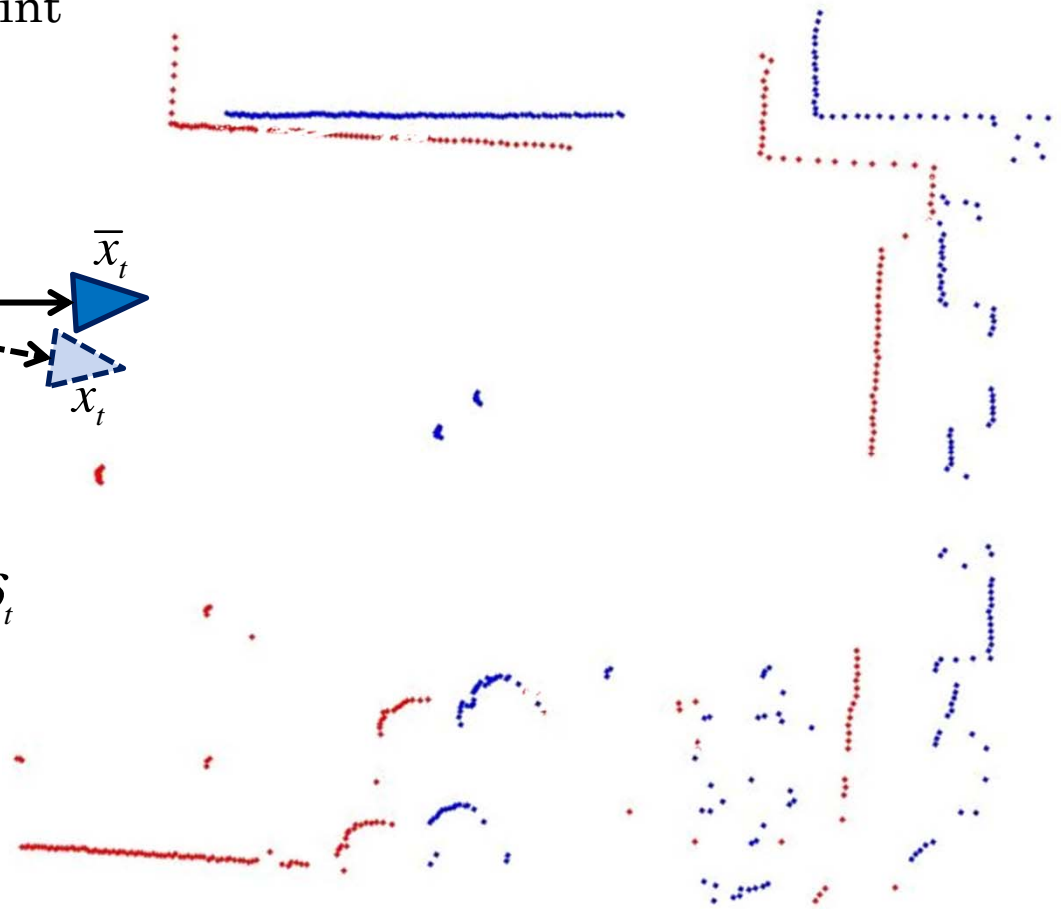
GRAPHSLAM WITH SCAN REGISTRATION

- True, odometry motion and resulting scans, can choose to include motion model constraint




$$x_t = g(x_{t-1}, u_t) + \delta_t$$

$$\delta_t \sim N(0, R_t)$$



GRAPHSLAM WITH SCAN REGISTRATION

- ICP scan match is a measurement between current and previous pose



$$y_t = h(x_t^r, x_{t-1}^r) + \delta_t$$

$$\delta_t \sim N(0, Q_t)$$

GRAPHSLAM WITH SCAN REGISTRATION

- Available information

- Inputs and Motion model

$$u_{0:t}, \quad x_t^r = g(x_{t-1}^r, u_t) + \delta_t$$

- Measurements and measurement model

$$y_{1:t}, \quad y_t = h(x_t^r, x_{t-1}^r) + \varepsilon_t = x_t^r - R_t^* x_{t-1}^r - t_t^*$$

- Where $y_t = 0$
- The scan registration process, therefore, changes the measurement model into a motion model
 - Depends on the current and previous robot state only
 - Can choose to include regular motion model too, and will be weighted based on relative uncertainty

GRAPHSLAM DERIVATION (LU/MILIOS)

- So the resulting negative log likelihood measurement constraint for each ICP match is

$$-\ln p(y_t | x_{t-1:t}) = \text{const.} + [y_t - h(x_{t-1:t})]^T Q_t^{-1} [y_t - h(x_{t-1:t})]$$

- In general, if loop closure is detected from scan i to scan j , we can add a measurement constraint between any two poses

$$-\ln p(y_{i,j} | x_{i,j}) = \text{const.} + [y_{i,j} - h(x_{i,j})]^T Q_{i,j}^{-1} [y_{i,j} - h(x_{i,j})]$$

- The full set of constraints collected are once again formed into a large optimization problem

GRAPHSLAM DERIVATION (LU/MILIOS)


- Again, we take the negative log likelihood version of the cost function

$$\min_{x_{0:t}} J = \text{const.} + \sum_{i,j} \left[y_{i,j} - h(x_{i,j}) \right]^T Q_{i,j}^{-1} \left[y_{i,j} - h(x_{i,j}) \right]$$



- Then solve the quadratic program however we'd like
 - Pseudo-inverse
 - Gauss-Newton
 - Levenberg-Marquardt

RESULTS FROM LU AND MILIOS

- Odometry only, and with GraphSLAM, using Sick Lidar [Lu, Milios at York in 1997].




GRAPHSLAM RESULTS




GRAPHSLAM RESULTS

- GPS/Odometry map of Stanford (600m x 600m)



GRAPHSLAM RESULTS

- Corrected using GraphSLAM




EXTENSIONS

- 3D GraphSLAM [Nuchter 2008]
 - Extension is actually taking the derivatives for linearization and moving it to 3D. Looked a lot at what the best way is for numerical stability of the solution when formulating the problem as a sequence of 3DOF inertial poses.
 - Euler angles
 - Quaternions
 - Helical motion
 - Rotation Linearization
 - ICP related improvements using KD-trees
 - Global Relaxation, a method for revisiting scan matching given the results of GraphSLAM

RESULTS FROM NUCHTER

- Large scale outdoor campus mapping



Odometry



Without loop closure

RESULTS FROM NUCHTER



With loop closure




With global relaxation

EXTENSIONS BY OLSON



- Re-parametrization [Olson 2009]
 - Most work uses a sequence of transformations between global poses to capture the motion of the robot
 - Olson uses an addition of differences in the pose parameters
 - This is inexact, but much faster
- Stochastic Gradient Descent
 - Do forever:
 - Pick a constraint
 - Descend in direction of constraint's gradient
 - Scale gradient magnitude by alpha/iteration
 - Clamp step size
 - iteration++
 - $\text{alpha/iteration} \rightarrow 0$ as $t \rightarrow \infty$
 - Robustness to local concavities
 - Hop around the state space, “stick” in the best one
 - Good solution very fast, “perfect” solution only as $t \rightarrow \infty$

RESULTS FROM OLSON

- Olson



Ground Truth





Noisy (simulated) input:

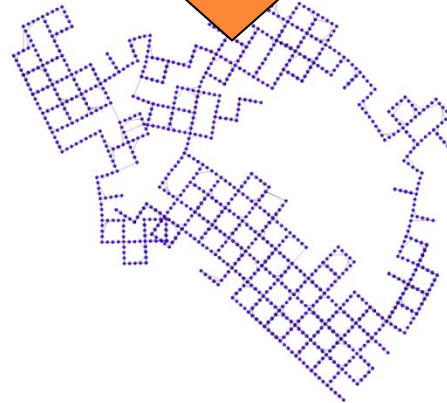

3500 poses

3499 temporal constraints



2100 spatial constraints



Gauss-Seidel, 60 sec.



*Multi-Level
Relaxation, 8.6 sec.*




Our method, 2.8 sec.

EXTENSIONS

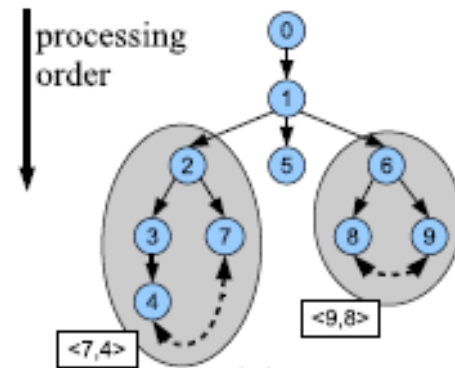

- Grisetti further modified the structure of the optimization by reorganizing the nodes of the graph into a tree with extra loop closing links.

[Grisetti 2010]

- A direct extension of Olson's formulation, using the




(b)



(c)

RESULTS FROM TORO

- 1000 nodes, less than a second to compute.



RESULTS FROM TORO

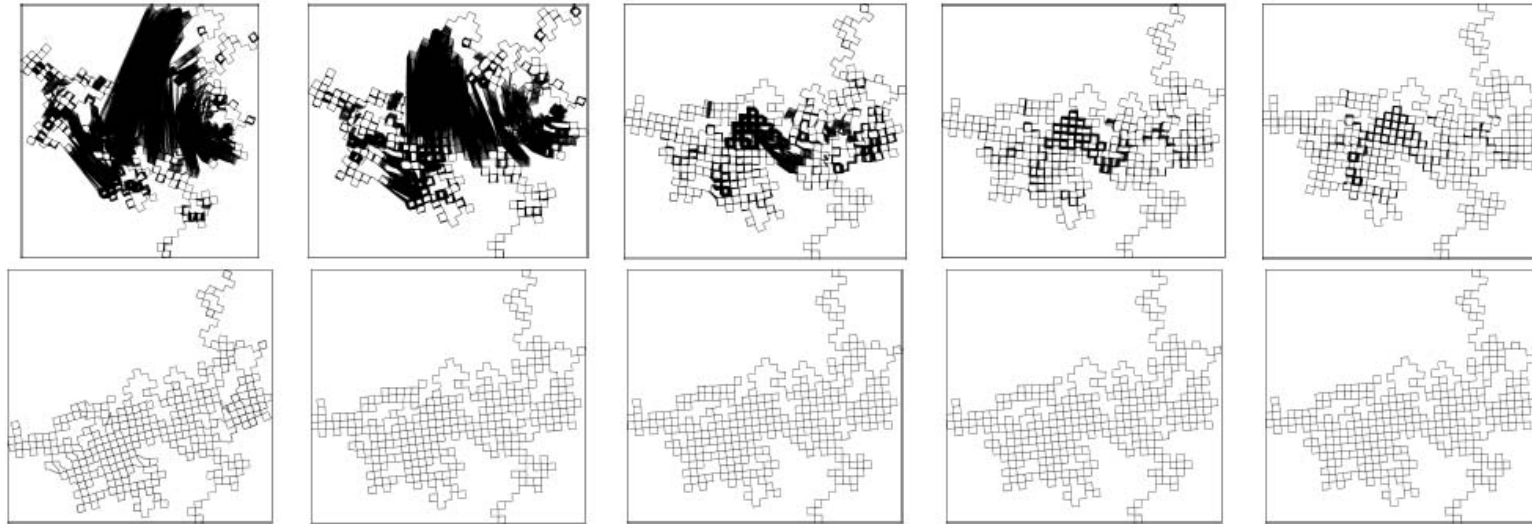




Fig. 4. Results of Olson's algorithm (first row) and our approach (second row) after 1, 10, 50, 100, 300 iterations for a network with 64k constraints. The black areas in the images result from constraints between nodes which are not perfectly corrected after the corresponding iteration (for timings see Figure 6).





RESULTS FROM TORO



Original

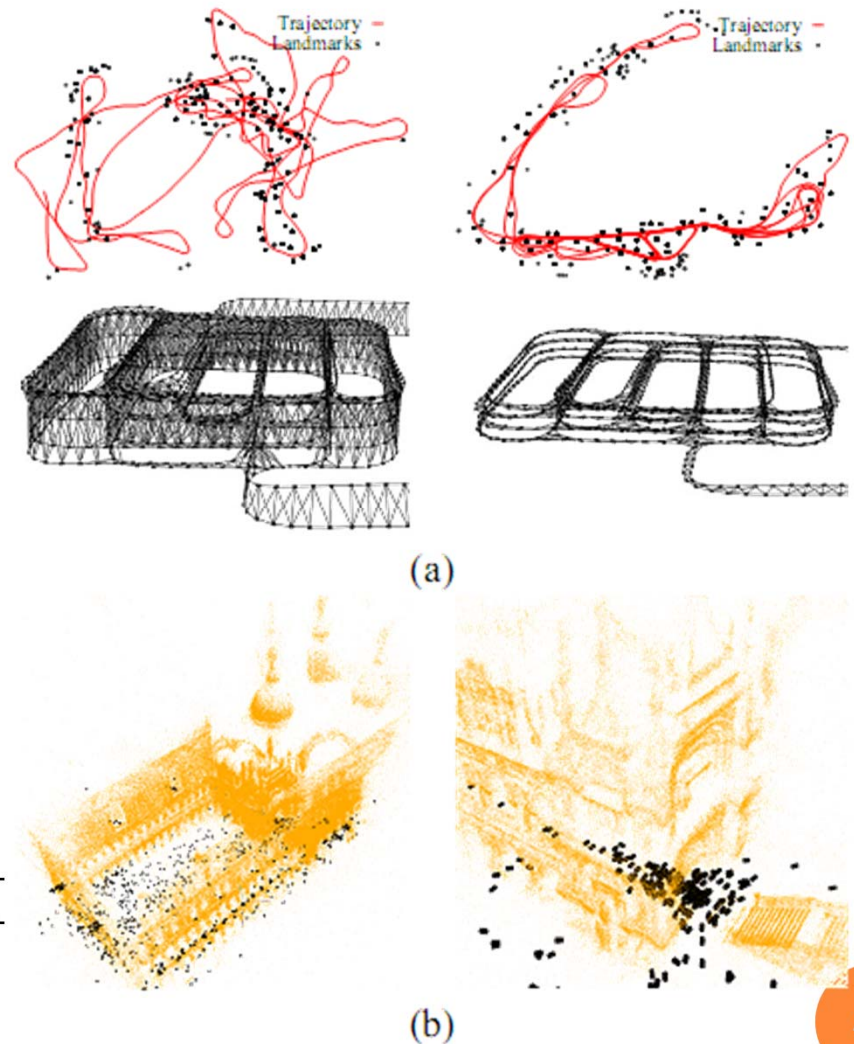


Optimized




CURRENT STANDARD – G2O

- Fast Backend Solver [Grisetti 2011]
- Takes the best of previous methods
- Works on wide range of problems
- Uses standard linear algebra packages
- Easily extensible, modifiable
- Available on OpenSLAM
- Integrated into ROS



RESULTS FROM NASA SAMPLE RETURN



STATE OF THE ART – [KINTINUOUS 2013]

Kintinuous 2.0

Real-time large scale dense loop closure with volumetric fusion mapping

Thomas Whelan*, Michael Kaess', John J. Leonard', John McDonald*

* Computer Science Department, NUI Maynooth

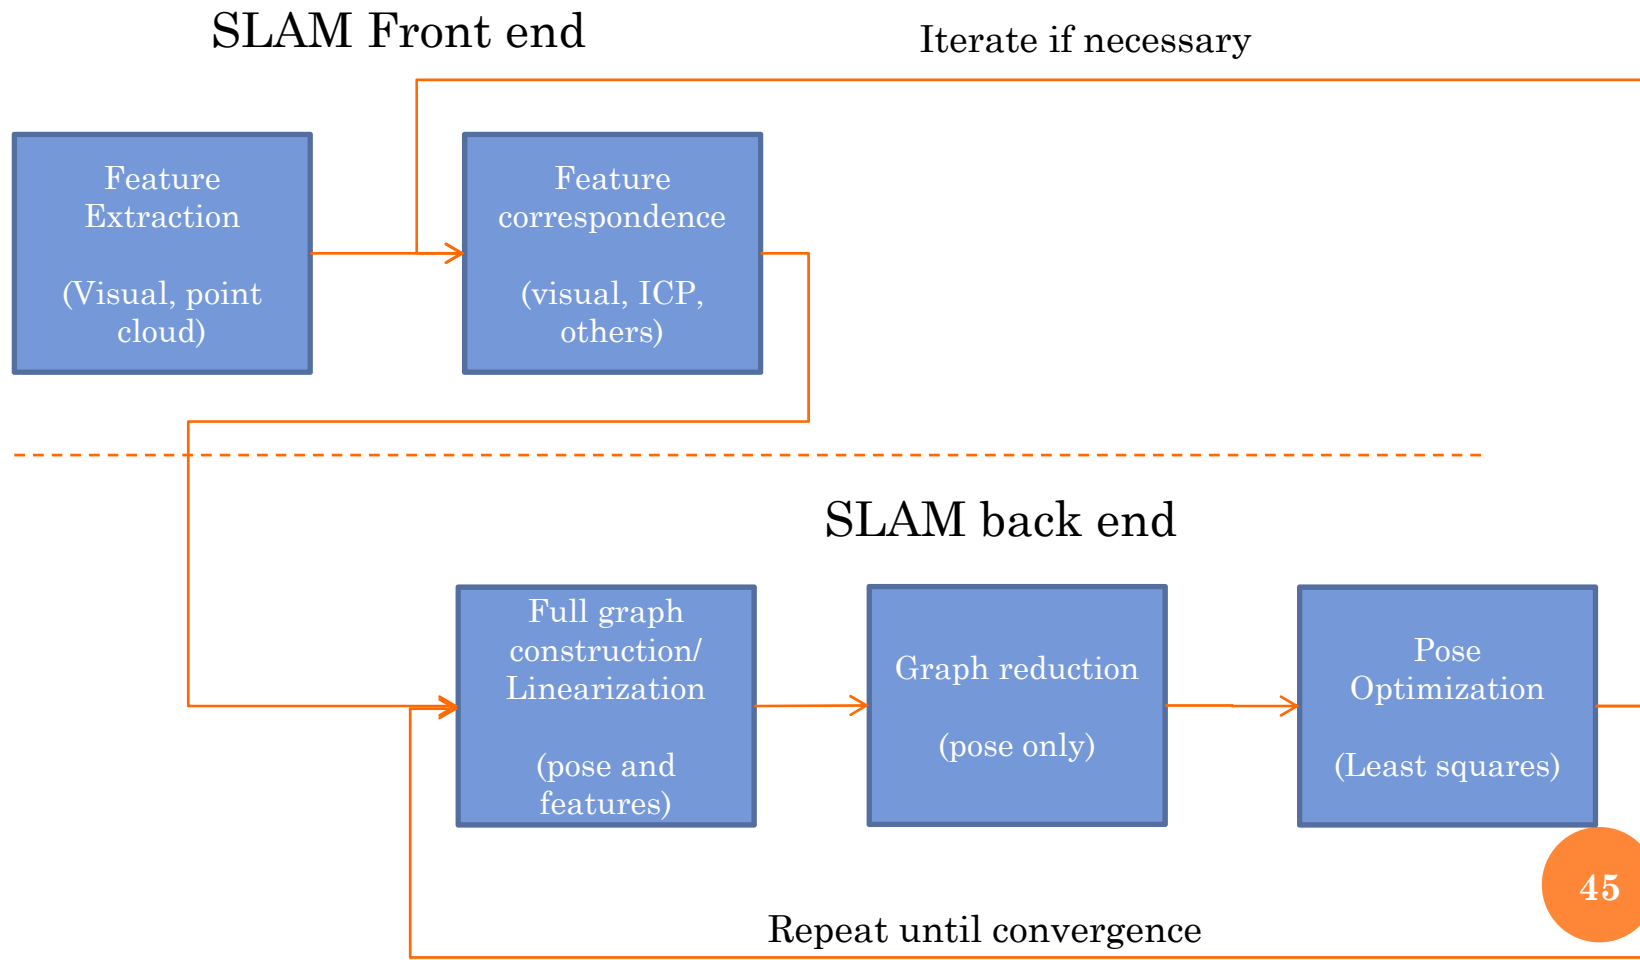
' Computer Science and Artificial Intelligence Laboratory, MIT

EXTRA SLIDES

GRAPHSLAM SOLUTION PIPELINE

- Feature extraction
 - Identify features in images (SIFT, SURF) or use laser scan points as features
- Feature correspondence
 - Standard visual techniques based on descriptors, proximity based such as ICP and many others
- Graph construction
 - Linearize measurement and motion information and populate a sparse matrix with constraint weightings based on covariance inverse (information matrix).
- Graph reduction
 - Reduce graph size by eliminating features, done by converting each feature measurement to an information gain on each pose from which it was measured
- Optimization
 - Any method you wish that can solve a sparse quadratic program (least squares, conjugate gradient, Levenberg-Marquardt)

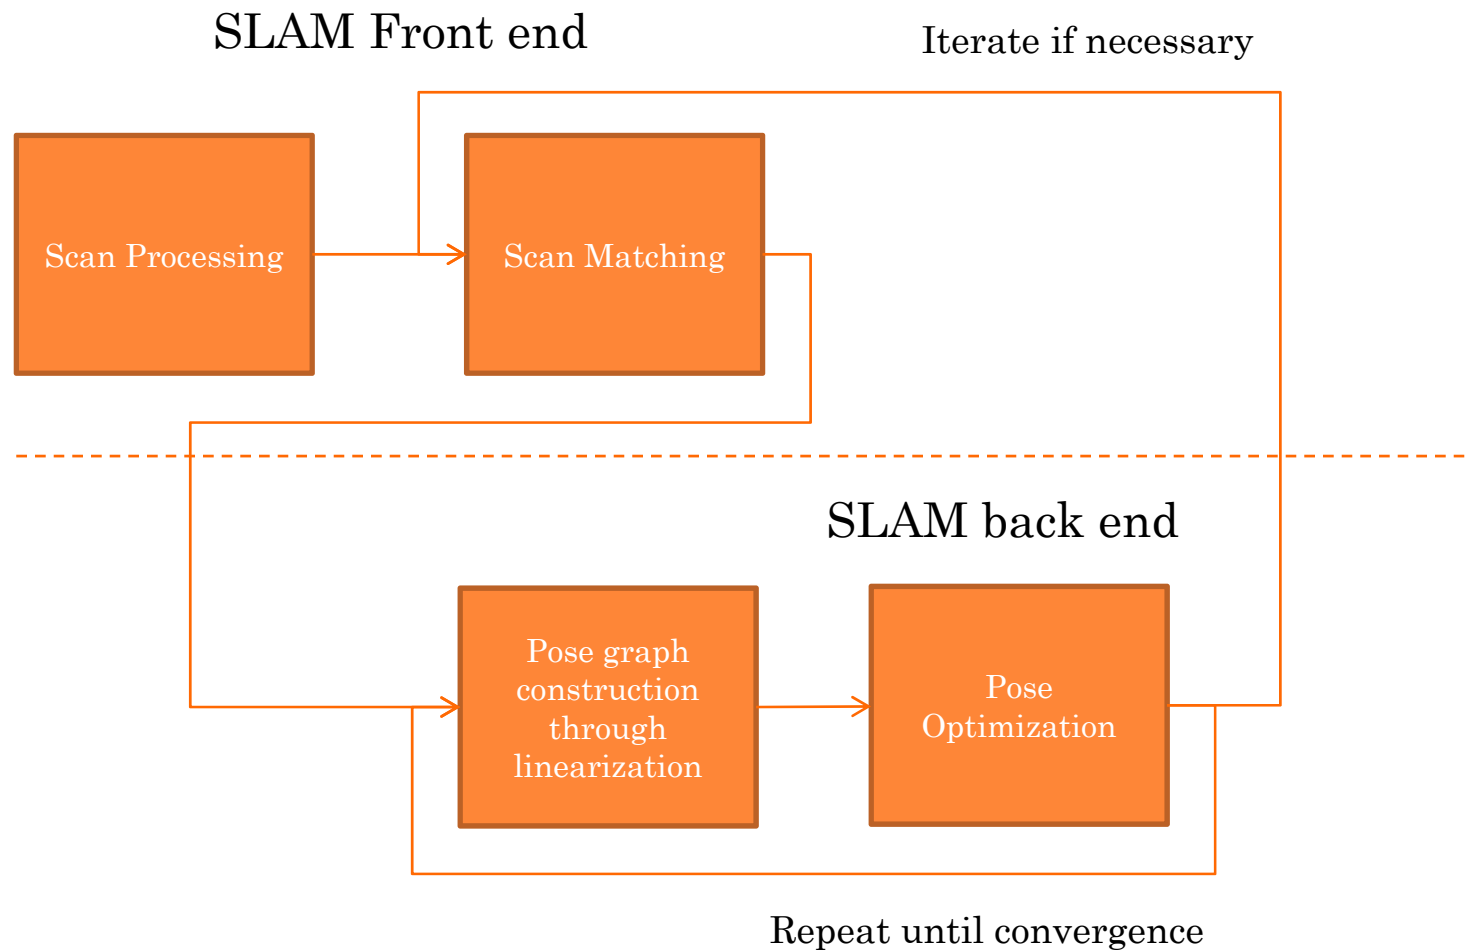
GRAPHSLAM OFFLINE SOLUTION PIPELINE



GRAPHSLAM DERIVATION (THRUN)

- There are four steps that result in Thrun's version of the GraphSLAM optimization
 1. Initialize: find an initial estimate of the trajectory, through odometry, raw ICP, whatever.
 2. Linearize: given the current estimate, find Jacobians of measurement and motion information, and construct the full graph.
 3. Reduce: eliminate the features from the graph through an explicit step, reducing graph size
 4. Solve: solve the quadratically constrained optimization to maximize probability of pose estimate, given measurements, inputs, correspondences.

GRAPHSLAM OFFLINE SOLUTION PIPELINE




GRAPHSLAM DERIVATION (THRUN)

- Derivation proceeds as follows:
 - Derivation of proposed optimization method
 - Define initial solution
 - Linearize measurement and motion models about current estimate of solution
 - Restate quadratic cost that results
 - Figure out how to reduce the graph to eliminate features, using marginals trick (Schur complement)
 - Present simple method for solving for the path and map
 - Repeat as necessary, updating initial solution each time through
 - Can add an outer loop that re-evaluates correspondence too

POPULATION OF SPARSE GRAPH MATRIX

- Motion constraint information
- Measurement constraint information



LINEARIZATION FOR COST DEFINITION

- A simple first order Taylor series expansion of $g(x,u)$ and $h(x,c)$ will result in a quadratic cost function
 - We can therefore proceed with sequential quadratic programming
- Similar to the EKF, the linearization is as follows

$$\begin{aligned} g(x_{t-1}, u_t) &\approx g(\mu_{t-1}, u_t) + \left. \frac{\partial}{\partial x_{t-1}} g(x_{t-1}, u_t) \right|_{x_{t-1}=\mu_{t-1}} (x_{t-1} - \mu_{t-1}) \\ &= g(\mu_{t-1}, u_t) + G_t \cdot (x_{t-1} - \mu_{t-1}) \end{aligned}$$

$$\begin{aligned} h(z_t, c_t^i) &\approx h(\mu_t, c_t^i) + \left. \frac{\partial}{\partial z_t} h(z_t, c_t^i) \right|_{z_t=\mu_t} (z_t - \mu_t) \\ &= h(\mu_t, c_t^i) + H_t \cdot (z_t - \mu_t) \end{aligned}$$

LINEARIZATION FOR MATRIX DEFINITION

- Substituting into the cost function gives

$$\begin{aligned} J = \text{const.} &+ \left[x_0 - \mu_0 \right]^T \Sigma_0^{-1} \left[x_0 - \mu_0 \right] \\ &+ \sum_t \left[x_t - g(\mu_{t-1}, u_t) - G_t \cdot (x_{t-1} - \mu_{t-1}) \right]^T R^{-1} \left[x_t - g(\mu_{t-1}, u_t) - G_t \cdot (x_{t-1} - \mu_{t-1}) \right] \\ &+ \sum_t \sum_i \left[y_t^i - h(\mu_t, c_t^i) - H_t \cdot (z_t - \mu_t) \right]^T Q^{-1} \left[y_t^i - h(\mu_t, c_t^i) - H_t \cdot (z_t - \mu_t) \right] \end{aligned}$$

- Which has many constant terms (mean is known) that can be combined into one.

LINEARIZED GRAPHSLAM OPTIMIZATION

- Rearranging, we can write the cost as

$$\begin{aligned} J = \text{const.} &+ x_0^T \Sigma_0^{-1} x_0 + \sum_t x_{t-1:t}^T \begin{pmatrix} -G_t^T \\ 1 \end{pmatrix} R^{-1} \begin{pmatrix} -G_t^T & 1 \end{pmatrix} x_{t-1:t} \\ &+ x_{t-1:t}^T \begin{pmatrix} -G_t^T \\ 1 \end{pmatrix} R^{-1} [g(\mu_{t-1}, u_t) - G_t \mu_{t-1}] \\ &+ \sum_t \sum_i z_t^T H_t^{iT} Q^{-1} H_t^i z_t + z_t^T H_t^{iT} Q^{-1} [y_t^i - h(\mu_t, c_t^i) + H_t^i \mu_t] \end{aligned}$$

- Which is of the form, a simple least squares problem.

$$J = \text{const} - z_{0:t}^T \Omega z_{0:t} + z_{0:t}^T \xi$$

- Where the full information matrix and vector are the two coefficients in the cost function

CONSTRUCTING THE COST FUNCTION

- We need to construct both the information vector and matrix
 - All constraints can be added independently in negative log likelihood form, taking care to place the additions in the correct rows and column (per initial diagram)

- Prior $\Omega = \Omega_0$

- Each motion step
$$\Omega = \Omega + \begin{pmatrix} -G_t^T \\ 1 \end{pmatrix} R^{-1} \begin{pmatrix} -G_t^T & 1 \end{pmatrix}$$

$$\xi = \xi + \begin{pmatrix} -G_t^T \\ 1 \end{pmatrix} R^{-1} [g(\mu_{t-1}, u_t) - G_t \mu_{t-1}]$$

- Each measurement
$$\Omega = \Omega + H_t^{iT} Q^{-1} H_t^i$$
$$\xi = \xi + H_t^{iT} Q^{-1} [y_t^i - h(\mu_t, c_t^i) + H_t^i \mu_t]$$

POPULATION OF SPARSE GRAPH MATRIX

- So, after linearization, we can form a quadratic cost matrix in all of the decision variables, which looks like

$$\Omega = \begin{array}{c} \begin{array}{cc} & \begin{array}{cccc} & x_{0:t} & & m \end{array} \\ \begin{array}{c} x_{0:t} \\ \\ \\ \\ \\ \\ \\ \\ \end{array} \left[\begin{array}{cccc|cccc} \cdot & \cdot & & & \cdot & \cdot & & & \\ \cdot & \cdot & \cdot & & \cdot & & \cdot & & \\ & \cdot & \cdot & \cdot & & & & & \\ & & \cdot & \cdot & \cdot & & & & \\ & & & \cdot & \cdot & \cdot & & & \\ & & & & \cdot & \cdot & \cdot & & \\ & & & & & \cdot & \cdot & & \\ & & & & & & \cdot & \cdot & \\ & & & & & & & \cdot & \cdot \end{array} \right] \\ \begin{array}{c} m \\ \\ \\ \\ \\ \\ \\ \\ \end{array} \left[\begin{array}{cccc|cccc} \cdot & & & & \cdot & & & & \\ \cdot & & & \cdot & \cdot & & & & \cdot & \\ & \cdot & & & & & & & & \\ & & \cdot & \cdot & & & & & & \\ & & & \cdot & \cdot & & & & & \\ & & & & \cdot & & & & & \\ & & & & & \cdot & \cdot & & & \\ & & & & & & \cdot & \cdot & & \\ & & & & & & & \cdot & \cdot & \end{array} \right] \end{array}$$

GRAPH REDUCTION

- The next big step is to eliminate all the features from the cost formulation, to reduce the size of the optimization problem
- The full SLAM posterior can be factored

$$p(z_{0:t} | y_{1:t}, u_{1:t}, c_{1:t}) = p(x_{0:t} | y_{1:t}, u_{1:t}, c_{1:t}) p(m | x_{0:t}, y_{1:t}, u_{1:t}, c_{1:t})$$

- Where we can find the marginal pose distribution by integrating over map variables

$$p(x_{0:t} | y_{1:t}, u_{1:t}, c_{1:t}) = \int p(z_{0:t} | y_{1:t}, u_{1:t}, c_{1:t}) dm$$

GRAPH REDUCTION

- To find this marginal probability, we need a famous lemma:
 - Marginals of a multivariate distribution (Schur complement, inversion lemma):

Let the probability distribution $p(x, y)$ over the random variables x, y be a Gaussian represented in the information form:

$$\Omega = \begin{bmatrix} \Omega_{xx} & \Omega_{xy} \\ \Omega_{yx} & \Omega_{yy} \end{bmatrix} \text{ and } \xi = \begin{bmatrix} \xi_x \\ \xi_y \end{bmatrix}.$$

If Ω_{yy} is invertible, the marginal $p(x)$ is a Gaussian whose information form is

$$\bar{\Omega}_{xx} = \Omega_{xx} - \Omega_{xy} \Omega_{yy}^{-1} \Omega_{yx} \text{ and } \bar{\xi}_x = \xi_x - \Omega_{xy} \Omega_{yy}^{-1} \xi_y$$

GRAPH REDUCTION

- The elimination of features proceeds by using this lemma

- Applying the marginalization lemma

$$\Omega = \begin{bmatrix} \Omega_{x_{0:t}x_{0:t}} & \Omega_{x_{0:t}m} \\ \Omega_{mx_{0:t}} & \Omega_{mm} \end{bmatrix} \text{ and } \xi = \begin{bmatrix} \xi_{x_{0:t}} \\ \xi_m \end{bmatrix}$$

$$\bar{\Omega}_{x_{0:t}x_{0:t}} = \Omega_{x_{0:t}x_{0:t}} - \Omega_{x_{0:t}m} \Omega_{mm}^{-1} \Omega_{mx_{0:t}}$$

$$\bar{\xi}_{x_{0:t}} = \xi_{x_{0:t}} - \Omega_{x_{0:t}m} \Omega_{mm}^{-1} \xi_m$$

$$\bar{\Omega}_{x_{0:t}x_{0:t}} = \Omega_{x_{0:t}x_{0:t}} - \sum_i \Omega_{x_{0:t}m_i} \Omega_{m_i m_i}^{-1} \Omega_{m_i x_{0:t}}$$

$$\bar{\xi}_{x_{0:t}} = \xi_{x_{0:t}} - \sum_i \Omega_{x_{0:t}m_i} \Omega_{m_i m_i}^{-1} \xi_{m_i}$$

The matrix $\Omega_{x_{0:t}m_i}$ is nonzero only for poses in which the feature was measured.

- But this seems to require an inversion the size of the map features
- Luckily, each feature is independent of all other features, so it is a block diagonal inversion, that can be done one feature at a time

SOLVING THE REDUCED OPTIMIZATION

- Finally, we solve the reduced quadratic program by simply inverting the information matrix, and recover the robot poses

$$\begin{aligned}\bar{\Sigma}_{x_{0:t}x_{0:t}} &= \bar{\Omega}_{x_{0:t}x_{0:t}}^{-1} \\ \bar{\mu}_{x_{0:t}} &= \bar{\Sigma}_{x_{0:t}x_{0:t}} \bar{\xi}_{x_{0:t}}\end{aligned}$$

- This inversion would be very fast if every feature was observed at only one time, but is actually a little dense due to loop closures. We have choices, but all must do some serious work to get a solution.
 - Pseudo-inverse
 - Gauss-Newton
 - Levenberg-Marquardt

SOLVING THE REDUCED OPTIMIZATION

- To recover the map (if needed), we want to solve for the conditional map probability

$$p(m \mid x_{0:t}, y_{1:t}, u_{1:t}, c_{1:t})$$

- To find this conditional probability, we need the conditioning lemma

Let the probability distribution $p(x, y)$ over the random variables x, y be a Gaussian represented in the information form:

$$\Omega = \begin{bmatrix} \Omega_{xx} & \Omega_{xy} \\ \Omega_{yx} & \Omega_{yy} \end{bmatrix} \text{ and } \xi = \begin{bmatrix} \xi_x \\ \xi_y \end{bmatrix}.$$

The conditional $p(x \mid y)$ is a Gaussian whose information matrix is Ω_{xx} and whose information vector is $\xi_x - \Omega_{xy} y$

RECOVERING THE MAP

- Application of the lemma yields

$$\begin{aligned}\bar{\Sigma}_{mm} &= \Omega_{mm}^{-1} \\ \bar{\mu}_m &= \Sigma_{mm} \left(\xi_m - \Omega_{mx_{0:t}} \bar{\mu}_{x_{0:t}} \right)\end{aligned}$$

- The feature locations are finally computed by again applying the fact that each feature is independent, so that we get a simple feature by feature reconstruction
 - Define the set of poses at which feature i was observed as $\tau(i)$

$$\begin{aligned}\Sigma_{m_i m_i} &= \Omega_{m_i m_i}^{-1} \\ \mu_{m_i} &= \Sigma_{m_i m_i} \left(\xi_{m_i} - \Omega_{m_i \tau(i)} \bar{\mu}_{\tau(i)} \right)\end{aligned}$$

NUCHTER'S OPENSAM RESULTS

