

Efficient Information-Theoretic Graph Pruning for Graph-Based SLAM with Laser Range Finders

Henrik Kretzschmar

Cyrill Stachniss

Giorgio Grisetti

Abstract—In graph-based SLAM, the pose graph encodes the poses of the robot during data acquisition as well as spatial constraints between them. The size of the pose graph has a substantial influence on the runtime and the memory requirements of a SLAM system, which hinders long-term mapping. In this paper, we address the problem of efficient information-theoretic compression of pose graphs. Our approach estimates the expected information gain of laser measurements with respect to the resulting occupancy grid map. It allows for restricting the size of the pose graph depending on the information that the robot acquires about the environment or based on a given memory limit, which results in an any-space SLAM system. When discarding laser scans, our approach marginalizes out the corresponding pose nodes from the graph. To avoid a densely connected pose graph, which would result from exact marginalization, we propose an approximation to marginalization that is based on local Chow-Liu trees and maintains a sparse graph. Real world experiments suggest that our approach effectively reduces the growth of the pose graph while minimizing the loss of information in the resulting grid map.

I. INTRODUCTION

Maps of the environment are needed for a wide range of robotic applications. In the past, several effective approaches to mapping have been developed. The graph-based formulation of the simultaneous localization and mapping (SLAM) problem models the poses of the robot as nodes in a graph. Spatial constraints between poses resulting from observations and odometry are encoded as edges. Graph-based approaches typically marginalize out features or local grid maps and reduce the mapping problem to trajectory estimation without prior map knowledge.

Most of the SLAM approaches assume that map learning is carried out as a preprocessing step and that the robot then uses the acquired model for tasks such as localization and path planning. A robot that has to extend the map of its environment during long-term operation cannot apply most of the existing graph-based mapping approaches since their complexity grows with the length of the robot’s trajectory. The reason for this is that standard graph-based approaches constantly add new nodes to the graph. As a result, memory and computational requirements grow over time, preventing long-term mapping applications. A constantly growing graph slows down graph optimization and makes it more and more



Fig. 1: The goal of our work is to prune the SLAM pose graph (top), maintaining a sparse pose graph (bottom), while minimizing the loss of information in the graph and the resulting map.

costly to find constraints between the current pose and former poses, i.e., to identify loop closures.

In this paper, we present a novel information-theoretic pruning approach that allows graph-based SLAM systems to operate in static environments over extended periods of time. Fig. 1 depicts a motivating example. The top image shows the pose graph and the resulting map obtained by a standard graph-based approach to SLAM. The bottom image displays the corresponding pose graph along with the map resulting from our information-theoretic pruning approach.

This paper makes two contributions. First, we present an approach to select laser scans for removal such that the expected loss of information with respect to the map is minimized. Our unbiased selection applies the information-theoretic concept of mutual information. The second contribution is an efficient approximation to marginalize out the pose nodes corresponding to the discarded laser scans. Marginalizing out a pose node from the graph means summarizing the information in the edges that are incident to that node in the edges between the nodes that are kept. The fill-in caused by exact marginalization, however, causes the pose graph to become dense and therefore leads to high memory usage and furthermore dramatically slows down graph optimization. In contrast to that, our approach preserves the sparse structure of the graph while summarizing most of the information in the constraints between poses that are kept. To maintain sparsity, our method uses Chow-Liu trees to locally approximate elimination cliques. The combination of both techniques is highly relevant to long-term mapping, particularly when the robot frequently re-traverses already visited areas. Our approach also allows us to build an any-space SLAM system that aims at minimizing the expected loss of information.

All authors are with the University of Freiburg, Department of Computer Science, Freiburg, Germany. G. Grisetti is also with Sapienza University of Rome, Department of Systems and Computer Science, Rome, Italy.

This work has partly been supported by the DFG under SFB/TR-8, by the EC under FP7-231888-EUROPA, and by Microsoft Research, Redmond. The authors would like to thank Maximilian Beinhofer for fruitful discussions and Dirk Hähnel for providing the FHW and the Intel Research Lab datasets.

II. RELATED WORK

Most of the current graph-based approaches to SLAM do not provide means to effectively prune the pose graph. Instead, they add more and more nodes to the graph over time. One way to limit the number of nodes in the graph is to sample the trajectory of the robot at an appropriate spatial decimation [5]. A similar method is to only add a new node to the graph if it is not spatially close to any existing node [9]. Konolige and Bowman [10] presented an approach to lifelong mapping that uses a single stereo camera and that is able to update the map when the environment changes. Their method discards views based on a least-recently used algorithm. The above-mentioned techniques do not rely on information-theoretic concepts to determine which measurements to discard.

In contrast to that, Davison [3] analyzes mutual information, particularly in the case of Gaussian probability distributions, to guide image processing. In the vision community, Snavely *et al.* [15] aim to find a skeletal subgraph with the minimum number of interior nodes that spans the full graph while achieving a bound on the full covariance. Their technique is used for reconstructing scenes based on large, redundant photo collections. Kaess and Dellaert [8] examine the information that measurements contribute to the state estimate in the iSAM framework. In contrast to that, our approach estimates the mutual information of laser scans and the occupancy grid map. Ila *et al.* [7] propose to only incorporate non-redundant poses and informative constraints based on the relative distance between poses in information space and the expected information gain of candidate loop closures. As opposed to our maximum-likelihood approach to SLAM based on pose graphs, their method applies an information filter and does not marginalize out already added poses.

Recently, Eade *et al.* [4] presented a view-based monocular SLAM system that reduces the complexity of the graph by marginalization and subsequent suppression of edges incident to nodes of high degrees. Their heuristic discards the constraints that most agree with the current state estimate. This, however, introduces a bias into the system.

Unlike our previous work [12], which neglects the uncertainty of constraints, our approach presented in this paper uses Chow-Liu trees to efficiently preserve sparsity while minimizing the loss of information without biasing the system.

III. GRAPH-BASED SLAM

Graph-based approaches to SLAM model the poses of the robot as nodes in a graph. The edges of the graph model spatial constraints between the nodes. These constraints naturally arise from odometry measurements and from feature observations or scan matching. The so-called SLAM front-end interprets the sensor data to extract the spatial constraints. The so-called SLAM back-end typically applies optimization techniques to estimate the configuration of the nodes that best matches the spatial constraints.

A. The SLAM Front-end

Our laser-based front-end uses correlative scan matching to estimate a constraint between the current node and the

previous node. Our method also generates loop closure hypotheses by matching the current laser scan against a set of scans that is determined by the relative positional uncertainties and then rejects false hypotheses using the spectral clustering approach described by Olson [14]. Our method incrementally optimizes the pose graph while adding the poses and the constraints to it. Once the poses are estimated, the laser scans are used to render an occupancy grid map of the environment. The robot therefore needs to store the laser scans corresponding to the pose nodes.

B. The SLAM Back-end

The back-end aims at finding the spatial configuration \mathbf{x}^* of the nodes that maximizes the log likelihood of the observations. Let $\mathbf{x} = (x_1^T, \dots, x_n^T)^T$ be a vector where x_i describes the pose of node i . Let z_{ij} and Ω_{ij} be the mean and the information matrix of an observation of node j seen from node i assuming Gaussian noise. Let $\mathbf{e}_{ij}(\mathbf{x})$ be an error vector which expresses the difference between an observation and the current configuration of the nodes. Let \mathcal{C} be the set of pairs of nodes for which a constraint exists. Assuming the constraints to be independent, we have

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmin}} \sum_{\langle i,j \rangle \in \mathcal{C}} \mathbf{e}_{ij}(\mathbf{x})^T \Omega_{ij} \mathbf{e}_{ij}(\mathbf{x}). \quad (1)$$

Our approach applies the technique proposed in [6], which uses sparse Cholesky factorization to efficiently solve the system of linearized equations that is obtained from Eq. (1).

C. Pose Graphs as Gaussian Markov Random Fields

A pose graph can be seen as a Gaussian Markov random field (GMRF) that models the belief of the robot. Markov random fields are useful to analyze the conditional independence structure of probability distributions. GMRFs are a subset of Markov random fields suitable for expressing multivariate Gaussian distributions. Thus, a pose graph can directly be interpreted as a GMRF that is given by a graph $G = (V, \Psi)$. The nodes in the pose graph correspond to the random variables in V . All the potential functions in Ψ are either unary or binary. The unary potentials are typically set to unity, except the one for the initial pose, which is bound to the origin. The likelihood functions associated with the edges of the pose graph correspond to the binary potentials in Ψ . The belief is given by

$$p(x_1, \dots, x_t) = \frac{1}{Z} \prod_{i=1}^t \psi_i(x_i) \prod_{\langle i,j \rangle \in \mathcal{C}} \psi_{i,j}(x_i, x_j). \quad (2)$$

IV. SELECTING THE MOST INFORMATIVE LASER SCANS

Our first contribution is an approach to select the laser scans that are most informative with respect to the map estimate. Our technique aims at minimizing the expected loss of information in the resulting map without introducing a bias during the selection of the laser scans. Such a technique is important to allow for long-term robot mapping since a robot that keeps all scans will run out of resources at some point. In addition to that, our method can be used to

directly implement an any-space SLAM system. Whenever the memory limit is reached, our algorithm discards the laser scans that are expected to be least informative about the map and marginalizes out the corresponding pose nodes.

A. Finding the Most Informative Subset of Laser Scans

We define the map M as a random variable describing the state of the world. It is highly correlated to the random variables $Z_{1:t}$ describing the laser scans $z_{1:t}$ recorded at the poses $x_{1:t}$. We use Z_i^j to refer to an individual beam of laser scan Z_i . To estimate the state of the world m , we consider the posterior probability distribution of the map M given the laser measurements $z_{1:t}$. In this section, we are interested in finding the subset $Z^* \subseteq Z_{1:t}$ of at most n laser measurements that is expected to result in the smallest uncertainty about the map M .

Following the notation of [13], the average reduction in the uncertainty of the map M due to a set Z of laser measurements is given by the mutual information

$$I(M; Z) = H(M) - H(M | Z). \quad (3)$$

Hence, we want to find the subset $Z^* \subseteq Z_{1:t}$ of at most n laser measurements such that the mutual information of the map M and the subset Z^* is maximized, giving

$$Z^* = \operatorname{argmax}_{Z \subseteq Z_{1:t}, |Z| \leq n} H(M) - H(M | Z). \quad (4)$$

The conditional entropy $H(M | Z)$ of the map M given the set Z of measurements is the expected value, over the space of all possible measurements, of the conditional entropy of the map given the individual measurements z :

$$H(M | Z) = \int_z p(z) H(M | Z = z) dz. \quad (5)$$

B. Efficiently Estimating Mutual Information

Unfortunately, computing the conditional entropy given in Eq. (5) is infeasible without approximations since integrating over the space of all possible combinations of up to n laser measurements is practically impossible. In addition to that, computing the entropy $H(M | Z = z)$ of a map given a set of measurements z typically requires model assumptions about the world.

To efficiently compute $H(M | Z)$, we make the following assumptions. We assume the laser measurements and the individual laser beams to be independent. Furthermore, we model the map M as an occupancy grid map, i.e., a grid of independent discrete binary random variables C that take the values $\text{Val}(C) = \{\text{"free"}, \text{"occupied"}\}$. The entropy of an occupancy grid map M given a set of measurements z is then given by

$$\begin{aligned} H(M | Z = z) &= \sum_{C \in M} H(C | Z = z) = \\ &= - \sum_{C \in M} \sum_{c \in \text{Val}(C)} P(C = c | z) \log P(C = c | z). \end{aligned} \quad (6)$$

In addition to that, we ignore the distribution over $x_{1:t}$ and operate on the most likely estimate $x_{1:t}^*$, which is given

in Eq. (1). Furthermore, similar to most works on robot localization, we assume the likelihood of sensing objects to decrease with range. The a-priori probability of the j -th beam of a range measurement z_i , denoted as z_i^j , without any knowledge of the map M can be described by the exponential distribution

$$p(z_i^j) = \begin{cases} \eta \lambda e^{-\lambda z_i^j} & z_i^j \leq z_{\max}, \\ 0 & z_i^j > z_{\max}, \end{cases} \quad (7)$$

where z_{\max} is the maximum range of the scanner, λ a parameter of the measurement model, and η a normalizer.

There are three possible outcomes of a measurement of a laser beam with respect to a particular grid cell that is located along the ray of the beam and given no prior map information. The laser beam either traverses the cell and thus observes the cell as free, the laser beam ends in the cell and thus observes the cell as occupied, or the laser beam does not observe the cell. The probability distribution of the outcome can be computed by integrating over the density $p(z_i^j)$. For instance, the probability that the beam Z_i^j does not reach a particular grid cell C that is located along the ray of the beam is given by

$$P(Z_i^j \text{ does not observe } C) = \int_0^{d(x_i^*, C)} p(z_i^j) dz_i^j, \quad (8)$$

where $d(x_i^*, C)$ is the distance between the pose x_i^* (see Eq. (1)) from which the laser scan Z_i is taken and the border of the grid cell C . Similarly, we can compute the probability of the cell being observed as free and the probability of the cell being observed as occupied. To compute the mutual information

$$I(C; Z) = H(C) - \sum_{z \in \mathcal{A}_Z} P(z) H(C | Z = z) \quad (9)$$

of the grid cell C and the set Z of laser measurements, we need to consider the set \mathcal{A}_Z of all possible measurement outcomes z with respect to the grid cell C of all k laser scans that are recorded close enough to potentially measure C .

In general, the number of possible combinations of grid cell measurement outcomes is exponential in k . It is therefore practically infeasible to enumerate all the combinations in a tree. However, we use a standard sensor model, $p(c | z_i^j)$, for laser range scanners that assigns to each cell one of the three occupancy values l_{free} , l_{occ} , and l_0 . Since the effect of a set of observations on a particular cell does not depend on the *order* the measurements were obtained, this model allows us to efficiently prune the tree. In fact, the result only depends on the *number* of free and occupied observations, i.e., the histogram of measurement outcomes. Therefore, ignoring the order, the number of histograms that we have to compute is quadratic in k . Computing their probabilities is cubic in k . Fortunately, the number k of scans that the algorithm has to consider is typically bounded: First, the maximum measurement range of laser scanners restricts the set of scans that have to be considered. Second, our technique discards laser scans online while building the graph and thus k typically stays small during mapping. We can further reduce

the computational burden by only considering at most l laser scans when computing the histograms. One good way of choosing the l laser scans is selecting the ones with the highest likelihood of measuring the cell according to Eq. (7). This approximation yields a linear complexity in k .

Finally, the mutual information $I(M; Z)$ of the map M and the set Z of laser scans is given by

$$I(M; Z) = \sum_{C \in \mathcal{M}} I(C; Z). \quad (10)$$

All terms needed to compute Z^* in Eq. (4) are specified and can be computed or approximated efficiently.

C. Discarding Laser Scans Online

Our approach can be used in two ways. First, by introducing a bound on the total number of laser scans, our method results in an any-space SLAM system. Second, setting a threshold on the expected information gain of laser scans, our algorithm only keeps scans that are expected to provide at least a certain amount of information about the map.

Computing the subset Z^* of n laser measurements that most reduces the uncertainty about the map has been shown to be at least NP-hard [11]. Fortunately, the problem is submodular. Hence, greedily selecting measurements results in obtaining a set of measurements that is at most a constant factor (≈ 0.63) worse than the optimal set. Motivated by this insight, our approach estimates the subset Z^* by successively discarding laser scans. In each step, it discards the laser scan that is expected to be least informative.

V. APPROXIMATE MARGINALIZATION OF POSES

The second contribution of this paper is an efficient approximation to marginalize out the pose nodes that correspond to the laser scans that were discarded using the approach described in the previous section. Marginalizing out a pose node from the graph implies summarizing the information stored in the edges that connect that node to the edges between nodes that are kept. Our approach aims at preserving sparsity while, at the same time, it seeks to minimize the loss of information. We refer the reader to [1] for a discussion on how to compound and merge geometric constraints with associated errors.

A. Marginalization Ceases Sparsity

The GMRF framework helps us understand the effects of marginalization on the belief of the robot. Suppose the belief is given by the joint Gaussian probability distribution $p(\alpha, \beta)$ over two stacked vectors α and β which together comprise the poses of the robot. In information form, the belief $p(\alpha, \beta)$ can be expressed as

$$p(\alpha, \beta) = \mathcal{N}^{-1} \left(\begin{bmatrix} \xi_\alpha \\ \xi_\beta \end{bmatrix}, \begin{bmatrix} \Omega_{\alpha\alpha} & \Omega_{\alpha\beta} \\ \Omega_{\beta\alpha} & \Omega_{\beta\beta} \end{bmatrix} \right), \quad (11)$$

where ξ are information vectors and Ω information matrices. To remove β from the belief, we compute the marginal density $p(\alpha) = \int p(\alpha, \beta) d\beta$. In information form, marginalization




Fig. 2: Eliminating a pose from the robot’s belief. Left: GMRFs Right: Information matrices, where gray shades indicate nonzero values. Top: x_9 is slated for removal. Middle: Belief after marginalizing out node x_9 . The former neighbors of x_9 form an elimination clique making the graph dense. Bottom: Belief resulting from Chow-Liu tree approximation of the elimination clique.

requires computing the Schur complement over the variables which should be kept. Hence, the parameters of $p(\alpha)$ are

$$\xi = \xi_\alpha - \Omega_{\alpha\beta} \Omega_{\beta\beta}^{-1} \xi_\beta, \quad (12)$$

$$\Omega = \Omega_{\alpha\alpha} - \Omega_{\alpha\beta} \Omega_{\beta\beta}^{-1} \Omega_{\beta\alpha}. \quad (13)$$

Unfortunately, as, for instance, discussed in [5], the Schur complement in Eq. (13) introduces new constraints between all pairs of variables that are directly related to the eliminated variables, adding a so-called elimination clique to the graph. This fill-in destroys the natural sparsity pattern that is typical to SLAM problems. See Fig. 2 for an illustration.

Representing a dense matrix typically requires significantly more memory resources than representing a sparse matrix. Furthermore, the density adversely affects the computational costs of subsequent marginalizations and of the underlying pose graph optimization. In the worst case, the number of dependencies increases quadratically with the number of variables. Hence, exact marginalization ceases the sparsity of the belief and therefore introduces a complexity that is not suited for long-term map learning.

B. Approximate Marginalization

In this paper, we propose to reduce the number of constraints in the elimination cliques, which emerge when marginalizing out pose nodes. We motivate this technique by the fact that this part of the graph has just become dense due to the fill-in resulting from marginalization.

To measure the effect of the approximation, we make use of the relative entropy or Kullback-Leibler divergence, which is




Fig. 3: Chow-Liu tree approximation. Left: Graphical model of a probability distribution $p(\tilde{\mathbf{x}})$, where the mutual information of some pairs of variables is indicated by the numbers and by the thickness of the edges. Right: Chow-Liu tree approximation of $p(\tilde{\mathbf{x}})$.

a standard measure of the difference between two probability distributions p and q . It is given by

$$D_{KL}(p \parallel q) = \int_x p(x) \log \frac{p(x)}{q(x)} dx. \quad (14)$$

Let $\tilde{\mathbf{x}} = (x_1^T, \dots, x_k^T)^T$ be the stacked vector of the set of variables in the elimination clique. Furthermore, let $p(\tilde{\mathbf{x}})$ be the density which arises from the constraints within the clique. A key idea of our approach is to reduce the computational burden by approximating the density $p(\tilde{\mathbf{x}})$ with a distribution that has fewer conditional dependencies. This implies treating some pairs of variables as conditionally independent.

We propose to reduce the number of constraints in the elimination cliques by locally approximating the density $p(\tilde{\mathbf{x}})$ with a probability distribution $q(\tilde{\mathbf{x}})$ such that each variable is conditioned upon only one of the other variables:

$$p(\tilde{\mathbf{x}}) = p(x_k) \prod_{i=1}^{k-1} p(x_i | x_{i+1}, \dots, x_k) \quad (15)$$

$$\approx p(x_k) \prod_{i=1}^{k-1} p(x_i | x_{i+1}) = q(\tilde{\mathbf{x}}). \quad (16)$$

Consequently, the graphical model of $q(\tilde{\mathbf{x}})$ is tree-shaped. Note, however, that we do not propose to transform the entire graph into a tree, only the nodes of the resulting elimination cliques. Importantly, our approximation preserves the global graph structure and, in particular, global loop closures.

C. Chow-Liu Tree Approximation

Chow-Liu trees [2] approximate a probability distribution $p(\tilde{\mathbf{x}})$ by a distribution $q_{\text{opt}}(\tilde{\mathbf{x}})$ such that each variable is conditioned upon at most one other variable and such that the Kullback-Leibler divergence between $p(\tilde{\mathbf{x}})$ and $q_{\text{opt}}(\tilde{\mathbf{x}})$ is minimized. Let the mutual information graph of a probability distribution be a fully connected graph such that each edge between two nodes x_i and x_j has weight equal to the mutual information $I(x_i; x_j)$ of these variables. Chow and Liu proved that the optimal approximation $q_{\text{opt}}(\tilde{\mathbf{x}})$ with first-order tree dependence to the probability distribution $p(\tilde{\mathbf{x}})$ has the same structure as the maximum-weight spanning tree of the mutual information graph of $p(\tilde{\mathbf{x}})$. The maximum-weight spanning tree of a graph and thus the Chow-Liu tree approximation can be computed efficiently using Kruskal’s algorithm (see Fig. 3 for an illustration).

Davison [3] shows that in the Gaussian case the mutual information of two variables x_i and x_j can be efficiently




Fig. 4: In this long-term experiment, the total number of nodes was restricted to 200.




Fig. 5: In this long-term experiment, our algorithm discarded all laser scans whose expected information gain was below a threshold.

computed as

$$I(x_i; x_j) = \frac{1}{2} \log_2 \left(\frac{|\tilde{\Sigma}_{ii}|}{|\tilde{\Sigma}_{ii} - \tilde{\Sigma}_{ij} \tilde{\Sigma}_{jj}^{-1} \tilde{\Sigma}_{ji}|} \right), \quad (17)$$

where $\tilde{\Sigma}_{ij}$ refers to the entry in the covariance matrix $\tilde{\Sigma}$ of probability distribution $p(\tilde{\mathbf{x}})$ which relates x_i and x_j based on the constraints in the elimination clique. Consequently, our optimal tree-shaped approximation is local to the elimination clique. The covariance matrix $\tilde{\Sigma}$ can be obtained by inverting $p(\tilde{\mathbf{x}})$, i.e., the information matrix corresponding to the clique. This can be done efficiently since this matrix is bounded in size due to the sparsity of the graph.

In sum, our technique allows for efficient marginalization and preserves sparsity in the pose graph while seeking to minimize the loss of information.

VI. EXPERIMENTAL EVALUATION

To evaluate the presented approach, we carried out several experiments using a real ActivMedia Pioneer-2 robot equipped with a SICK laser range finder and applied it to benchmark datasets. We compare our pruning approach to the same SLAM system when no scans are discarded (referred to as “no pruning” or “standard approach”, see Sec. III).

A. Memory and Runtime Requirements

In this section, we analyze the memory requirements in terms of the size of the pose graph. In a first experiment, the robot moved around in our lab environment for an extended period of time (see Fig. 4 to Fig. 6). The plots clearly suggest that this setting leads to an explosion in terms of memory requirements when using the standard approach. This has a direct influence on the computational complexity for two reasons. First, the pose graph optimization scales with the number of edges, which grows roughly quadratically since

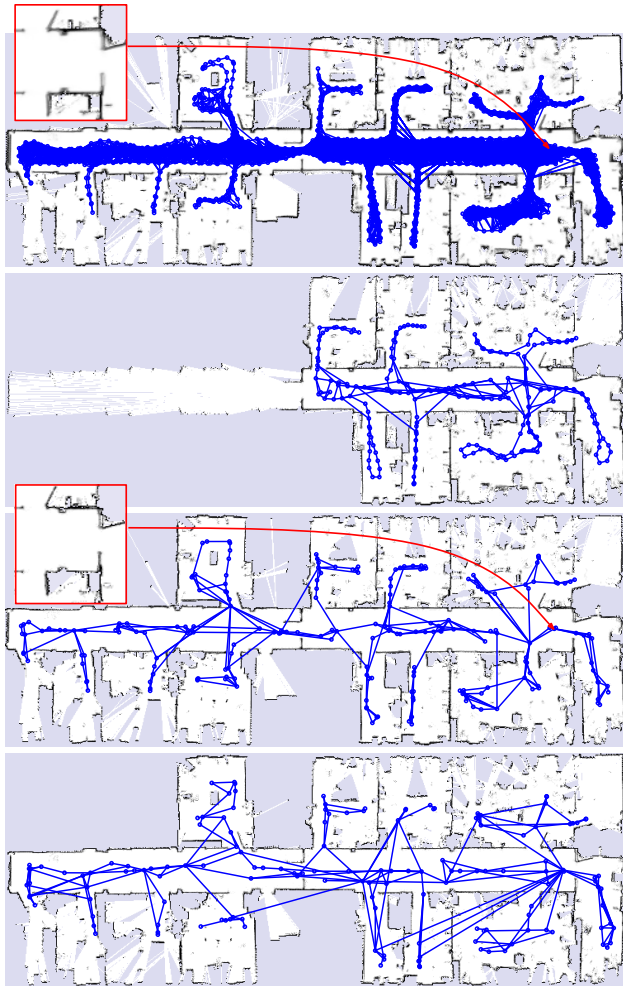


Fig. 6: Long-term experiment. First: Standard approach. 2597 laser scans, 15695 edges. Second: Our approach at an intermediate time step, 200 laser scans, 264 edges. Third: Our approach, 200 laser scans, 315 edges. Fourth: Our approach when setting a threshold on the mutual information, 148 laser scans, 250 edges.

the robot moves in the same environment. Second, the loop closing component of the SLAM front-end, which uses a scan matcher to find constraints between the current scan and all former scans that were recorded in the vicinity of the robot, has to consider an increasing number of nodes in each step. Our approach prunes the pose graph such that the number of nodes in the graph remains constant. Setting a threshold on the mutual information, the complexity does not grow as long as the robot does not explore new territory (see Fig. 5).

Fig. 6 depicts four pose graphs along with the corresponding maps. The first one depicts the pose graph of the standard approach. The second one shows the state of our approach before the robot entered the left side of the corridor. The limit of 200 nodes is used to model the right part only. The third image shows the pose graph modeling the entire environment. Note how our approach redistributed the nodes in the environment, still complying with the 200 node limit. Finally, the fourth image shows the map when setting a threshold on the mutual information. The motivating figure (Fig. 1) depicts another example.

Our approach saves computing time but also causes an overhead, which depends on the chosen parameters (particularly l , see Sec. IV-B) and on the environment that is mapped. This overhead is bounded since our algorithm constantly discards nodes. In our experiments, the speed of our pruning approach approximately ranges from running twice as fast as the standard approach to running four times slower than the standard approach. Our approach is beneficial when the robot frequently re-traverses already mapped areas. There is no gain if the robot mainly explores new territory.

B. Effects on the Most Likely Occupancy Grid Map

We furthermore analyzed the effects of our pruning technique on the resulting occupancy grid maps. We therefore compared the maps at a resolution of 10 cm and counted the number of cells that changed their most likely state (free, occupied, unknown) due to our pruning technique.

When mapping the Intel Research Lab, our pruning approach retained 349 of 1802 laser scans. As a consequence of this, 0.9% of the cells changed. In the long-term experiment, our method kept 148 of 2597 laser scans and 1.6% of the cells changed. When mapping the FHW, our approach maintained 250 of 2049 scans and 1.2% of the cells changed. Hence, the changes in the most likely maps are small.

C. Approximate Marginalization

Eliminating nodes from the belief by means of exact marginalization causes the pose graph to become dense. A sparse pose graph, however, is highly advantageous since it can be stored and optimized efficiently. Fig. 8 shows the evolution of the number of nodes and edges for both approaches while mapping the Intel Research Lab. Fig. 7 shows a comparison of the graphs resulting from exact marginalization and from our approximate marginalization technique when processing the Intel Research Lab dataset. The figure illustrates that our approach preserves the sparsity of the pose graph by using Chow-Liu trees to locally approximate the elimination cliques. On the right hand side, the figure depicts the 3σ covariance ellipses of the poses in the graphs. Our approach keeps less than 9% of the edges of the original graph (349 of 3916) but only 2.8% of the probability mass of the original pose graph is not covered by our approximation. The covariance estimates of our approach are typically more conservative (in this experiment by 41%) since less information is used during mapping.

D. Scan Alignment and Map Quality

In this section, we discuss how our online graph pruning technique may affect scan matching. Ignoring measurements leads to a belief with higher uncertainty. However, occupancy grid-based mapping approaches typically involve some form of scan alignment or scan matching to extract constraints. These systems have the following disadvantage when it comes to long-term map learning. Whenever the robot obtains a laser measurement, the scan matcher aims at aligning the new scan with existing scans in order to solve the data association problem. The probability that the scan matcher thereby makes




Fig. 7: Intel Research Lab. First: Standard approach, 1802 laser scans, 3916 edges. Arrows indicate misalignments in the map. Second: Exact marginalization of poses, 349 laser scans, 13052 edges. Third: Our approach preserves the sparsity of the pose graph, 250 laser scans, 349 edges. Fourth: 3σ covariance ellipses of the poses.




Fig. 8: Intel Research Lab. Left: The evolution of the number of nodes. Right: The evolution of the number of edges.

a small alignment error is nonzero. A scan that is incorporated at a slightly wrong position blurs the map. As a result, the probability that the scan matcher misaligns subsequent scans increases since scan matching is performed with misaligned scans. Hence, the probability of making alignment errors increases with the number of incorporated scans. In the long run, the map tends to become increasingly blurred and the mapping approach is likely to diverge. Fig. 6 and 7 depict the maps and graphs obtained from the Intel Research Lab dataset and the long-term experiments that we conducted in our office environment. The occupancy grid maps generated by the standard approach exhibit visibly more blur in several parts of the maps (see the arrows and the zoomed map view in the corresponding images). In general, the more often the robot re-traverses already visited terrain, the more blur is added to the maps. In contrast to the standard approach, our graph pruning method discards scans and thus produces maps with sharper obstacle boundaries even in cases in which the robot frequently re-traverses already visited places. Although we do not claim that these sharp maps are better estimates of the world, they reduce the risk of divergence in the mapping process.

VII. CONCLUSION

In this paper, we presented a method for efficient information-theoretic pruning of pose graphs in graph-based SLAM. Our approach seeks to select the most informative set of laser scans and allows for restricting the size of the pose graph either based on a memory limit, resulting

in an any-space mapping system, or based on a threshold on the minimum of information that a laser scan provides. When marginalizing out pose nodes in the graph, exact marginalization leads to a densely connected graph. To preserve the sparsity of the pose graph, we proposed to approximate the elimination cliques by local Chow-Liu trees, seeking to minimize the loss of information. Real world experiments illustrate the effectiveness of our method.

REFERENCES

- [1] P. Cheeseman and P. Smith. On the representation and estimation of spatial uncertainty. *Int. Journal of Robotics*, 5:56–68, 1986.
- [2] C. Chow and C. Liu. Approximating discrete probability distributions with dependence trees. *IEEE Transactions on Information Theory*, 14(3):462–467, 1968.
- [3] A.J. Davison. Active search for real-time vision. In *Proc. of the Int. Conf. on Computer Vision (ICCV)*, volume 1, 2005.
- [4] E. Eade, P. Fong, and M.E. Munich. Monocular graph SLAM with complexity reduction. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 3017–3024, Taipei, Taiwan, 2010.
- [5] R. Eustice, H. Singh, and J.J. Leonard. Exactly sparse delayed-state filters for view-based SLAM. *IEEE Transactions on Robotics*, 22(6):1100–1114, 2006.
- [6] G. Grisetti, R. Kümmerle, C. Stachniss, U. Frese, and C. Hertzberg. Hierarchical optimization on manifolds for online 2d and 3d mapping. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, Anchorage, AK, 2010.
- [7] V. Ila, J.M. Porta, and J. Andrade-Cetto. Information-based compact pose slam. *IEEE Transactions on Robotics*, 26(1):78–93, 2010.
- [8] M. Kaess and F. Dellaert. Covariance recovery from a square root information matrix for data association. *Journal of Robotics and Autonomous Systems (RAS)*, 57:1198–1210, Dec 2009.
- [9] K. Konolige and M. Agrawal. FrameSLAM: From bundle adjustment to realtime visual mapping. *IEEE Transactions on Robotics*, 24(5):1066–1077, 2008.
- [10] K. Konolige and J. Bowman. Towards lifelong visual maps. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 1156–1163, St. Louis, MO, USA, 2009.
- [11] A. Krause and C. Guestrin. Near-optimal nonmyopic value of information in graphical models. In *Proc. of Uncertainty in Artificial Intelligence (UAI)*, 2005.
- [12] H. Kretzschmar, G. Grisetti, and C. Stachniss. Lifelong map learning for graph-based SLAM in static environments. *KI – Künstliche Intelligenz*, 2010.
- [13] D.J.C. MacKay. *Information theory, inference, and learning algorithms*. Cambridge Univ Press, 2003.
- [14] E. Olson. *Robust and Efficient Robotic Mapping*. PhD thesis, MIT, Cambridge, MA, USA, June 2008.
- [15] N. Snavely, S.M. Seitz, and R. Szeliski. Skeletal graphs for efficient structure from motion. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, Anchorage, AK, 2008.