

Pose Graph Compression for Laser-Based SLAM

Cyrill Stachniss and Henrik Kretzschmar

Abstract The pose graph is a central data structure in graph-based SLAM approaches. It encodes the poses of the robot during data acquisition as well as spatial constraints between them. The size of the pose graph has a direct influence on the runtime and the memory requirements of a SLAM system since it is typically used to make data associations and within the optimization procedure. In this paper, we address the problem of efficient, information-theoretic compression of such pose graphs. The central question is which sensor measurements can be removed from the graph without losing too much information. Our approach estimates the expected information gain of laser measurements with respect to the resulting occupancy grid map. It allows us to restrict the size of the pose graph depending on the information that the robot acquires about the environment. Alternatively, we can enforce a maximum number of laser scans the robot is allowed to store, which results in an any-space SLAM system. Real world experiments suggest that our approach efficiently reduces the growth of the pose graph while minimizing the loss of information in the resulting grid map.

1 Introduction

Maps of the environment are needed for a wide range of robotic applications including transportation and delivery tasks, search and rescue, or efficient automated vacuum cleaning robots. The capability of building an appropriate model of the environment allows for designing robots that can efficiently operate in complex environments only based on their on-board sensors and without relying on external reference systems. In the past, several effective approaches to robot mapping have been developed. A popular approach to address the simultaneous localization and mapping (SLAM)

Cyrill Stachniss

University of Freiburg, Institute for Computer Science, e-mail: stachnis@informatik.uni-freiburg.de

Henrik Kretzschmar

University of Freiburg, Institute for Computer Science, e-mail: kretzsch@informatik.uni-freiburg.de




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

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. Often, graph-based approaches marginalize out features or local grid maps and reduce the problem to trajectory estimation without prior map knowledge, followed by mapping with known poses.

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 costly to find constraints between the current pose and former poses, i.e., to identify loop closures. There exist also incremental methods for online corrections that perform partial optimizations. These methods are mostly orthogonal to our contribution.

The contribution of this paper is an information-theoretic approach to lossy pose graph compression to allow graph-based SLAM systems to operate 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 compression approach. As can be seen, significantly less nodes are required to provide a comparable mapping result. 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 to determine the laser scans that should be removed. In order to keep the pose graph compact, the corresponding pose node needs to be eliminated from the pose graph. This is achieved by applying an approximate marginalization scheme that maintains the natural sparsity pattern that is observed in the context of SLAM. Our approach is highly relevant to long-term mapping, particularly when the robot frequently re-traverses already visited areas. It allows us to build an any-space SLAM system that aims at minimizing the expected loss of information.

2 Related Work

In the context of the SLAM problem, Lu and Milios [15] were the first who performed global map optimization using a pose graph. Since then, a large variety of approaches to minimize the error in pose graphs have been proposed [5, 7, 8, 11, 17, 19], as well as many others.

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. Some approaches group nodes into rigid local sub-maps [7] or subdivide the map into connected frames that contain maps that capture the local environment [1]. Typically, these methods do not discard nodes that store information about the environment and therefore do not prevent the graph from growing.

One way to reduce the number of nodes in the graph is to sample the trajectory of the robot at an appropriate spatial decimation [6]. A similar method is to only add a new node to the graph if it is not spatially close to any existing node [11]. Konolige and Bowman [12] 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. Similar to that, hierarchical techniques [8, 5, 17] have been employed to bound the computational requirements by optimizing only higher levels of the hierarchy.

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.* [20] 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.

The approach closest to our method is the work of Kaess and Dellaert [10]. Their work considers the information gain of measurements in the state estimate within the iSAM framework. In contrast to that, our approach estimates the mutual information of laser scans and the occupancy grid map, thus considering the effect on the resulting grid map explicitly. Ila *et al.* [9] 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 poses that were already added. 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.

When discarding laser range scans, our approach will marginalize out the corresponding pose node from the pose graph. Exact marginalization, however, would result in a dense pose graph and thus we apply an approximate marginalization

scheme [14] that is based on local Chow-Liu trees. Other robotics researchers also use Chow-Liu trees for approximating probability distributions, e.g., [2].

3 Brief Introduction to 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.

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 [18]. 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 stores the laser scans that correspond to the pose nodes in the pose graph.

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 , and 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. Furthermore, let $\mathbf{e}_{ij}(\mathbf{x})$ be an error vector which expresses the difference between an observation and the current configuration of the nodes and 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 [8], which uses sparse Cholesky factorization to efficiently solve the system of linearized equations that is obtained from Eq. (1).

4 Discarding Laser Scans by Information-theoretic Means

The main contribution of this paper 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.

4.1 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 MacKay [16], the average reduction in the uncertainty of the map M due to a set Z of laser measurements is given by the mutual information

$$I(M; Z) = H(M) - H(M | Z), \quad (2)$$

where H is the entropy. 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, i.e.,

$$Z^* = \underset{Z \subseteq Z_{1:t}, |Z| \leq n}{\operatorname{argmax}} H(M) - H(M | Z). \quad (3)$$

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 :

$$H(M | Z) = \int_z p(z) H(M | Z = z) dz \quad (4)$$

4.2 Efficiently Estimating Mutual Information

Unfortunately, computing the conditional entropy given in Eq. (4) 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 a standard 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

$$H(M | Z = z) = \sum_{C \in \mathcal{M}} H(C | Z = z) \quad (5)$$

$$= - \sum_{C \in \mathcal{M}} \sum_{c \in \text{Val}(C)} P(C = c | z) \log P(C = c | z). \quad (6)$$

To efficiently compute Z^* , we additionally 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 a specific object 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} denotes the maximum range of the scanner, λ is a parameter of the measurement model, and η is a normalizing constant.

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 the three cases, namely, the probabilities that (i) the beam Z_i^j does not reach a particular grid cell C that is located along the ray of the beam, (ii) the beam ends in that cell (measured as occupied), and (iii) the beam passes through that cell (measured as free) are given by

$$P(Z_i^j \text{ does not observe } C) = \int_0^{d_1(x_i^*, C)} p(z_i^j) dz_i^j \quad (8)$$

$$P(Z_i^j \text{ observes } C \text{ as occupied}) = \int_{d_1(x_i^*, C)}^{d_2(x_i^*, C)} p(z_i^j) dz_i^j \quad (9)$$

$$P(Z_i^j \text{ observes } C \text{ as free}) = \int_{d_2(x_i^*, C)}^{z_{\max}} p(z_i^j) dz_i^j, \quad (10)$$

where $d_1(x_i^*, C)$ is the distance between the pose x_i^* (see Eq. (1)) from which the laser scan Z_i is taken and the closest border of the grid cell C in the direction of the j^{th} beam (the border where the beam enters the cell). Similarly, $d_2(x_i^*, C)$ is the distance to the border of the grid cell C in the direction of the j^{th} beam that is furthest away from x_i^* (the border where the beam leaves the cell).

By exploiting Eq. (8) to (10), we can avoid integrating over all potential measurements as in Eq. (4). Instead, we can sum over all potential measurement outcomes.

This results in the mutual information

$$I(C; Z) = H(C) - \sum_{z' \in \mathcal{A}_Z} P(z') H(C | z') \quad (11)$$

of the grid cell C and the set Z of laser measurements. Here, we 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 the cell.

In general, the number of possible combinations of grid cell measurement outcomes is exponential in k as it is illustrated in the left image of Fig. 2. It is therefore practically infeasible to enumerate all the combinations in a tree. In our approach, we use a standard inverse measurement model, $p(c | z_i^j)$, for laser range scanners that updates each cell using one of the three 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* in which the measurements were obtained, this model allows us to efficiently combine nodes in the tree of all possible combinations. In fact, the result only depends on the *number* of free and occupied observations, i.e., the histogram of measurement outcomes, see Fig. 2 (right) for an illustration.




Fig. 2: Left: all combination of measurement combinations that can occur for n beams, here with $n = 2$. In the figure, f=free, o=occupied, and u=unknown. The number of combinations is 3^n . Right: Since the order of the measurements is irrelevant, the number of *possible outcomes* is quadratic in n .

By ignoring the order, the number of histograms that we have to compute is quadratic in k :

$$\#outcomes(k) = \sum_{i=0}^k \left(\sum_{j=0}^{k-i} 1 \right) = \sum_{i=0}^k (k-i+1) \quad (12)$$

$$= (k+1)^2 - \sum_{i=0}^k i = (k+1)^2 - \frac{(k+1)k}{2} \quad (13)$$

$$= \frac{k^2}{2} + \frac{3}{2}k + 1 \in \mathcal{O}(k^2) \quad (14)$$

If the individual probabilities of obtaining free/occupied/unknown measurements were equal for all laser scans, the probabilities of these outcomes could be computed by a multinomial distribution.

However, this is not the case here due to our model given in Eq. (7), which specifies that the likelihood of observing a cell depends on the distance from the view point to the cell under consideration. In our case, computing the probabilities of all outcomes requires cubic time in k . This can be achieved by using a hash table that is indexed by the number of free, occupied, and unknown measurements and that stores the accumulated probability mass for the corresponding outcome. By traversing the graph of possible outcomes (see right image of Fig. 2), from top to bottom, the computations can be performed as shown in Algorithm 1 (where $P(\langle \cdot, \cdot, \cdot \rangle)$ is implemented via a hash table) and the number of probabilities that need to be considered is

$$3 \sum_{i=1}^k \#outcomes(k) = \frac{3}{2}k^3 + \frac{9}{2}k^2 + 3k \in \mathcal{O}(k^3). \quad (15)$$

Thus, to compute the probability for each possible outcome is cubic in k , i.e. the number of measurements that can observe the grid cell C .

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 efficiently 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 C . This likelihood is given by $1 - P(Z_i^j \text{ does not observe } C)$, see Eq. (8). Thus, this approximation yields a linear complexity in k (for selecting the l scans out of k laser scans).

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

$$I(M; Z) = \sum_{C \in M} I(C; Z). \quad (16)$$

Algorithm 1 Compute probabilities for all measurement outcomes

Require: Z : set of laser measurements, C : cell

Ensure: $P(\langle \cdot, \cdot, \cdot \rangle)$: the probabilities of all outcomes (free, occupied, not observed)

$P(\langle 0, 0, 0 \rangle) = 1$

for $r = 1 \dots k$ **do**

for all $\langle f, o, u \rangle$ with $f + o + u == r$ **do**

$P(\langle f+1, o, u \rangle) += P(\langle f, o, u \rangle)P(Z_i^j \text{ observes } C \text{ as free});$

$P(\langle f, o+1, u \rangle) += P(\langle f, o, u \rangle)P(Z_i^j \text{ observes } C \text{ as occupied});$

$P(\langle f, o, u+1 \rangle) += P(\langle f, o, u \rangle)P(Z_i^j \text{ does not observe } C);$

end for

end for

return $P(\langle \cdot, \cdot, \cdot \rangle)$

All terms needed to compute Z^* in Eq. (3) are specified and can be computed or approximated efficiently.

4.3 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 for 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 [13]. 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.

5 Maintaining a Sparse Pose Graph

A pose graph can be seen as a Gaussian Markov random field (GMRF) that models the belief of the robot. In this view, each pose is a random variable that is represented as one node in the GMRF and each constraint between two poses in the pose graph is a binary potential between the nodes in the GMRF. Marginalizing out a pose node from the graph implies summarizing the information stored in the edges that connect that node in the edges between nodes that are kept. The main problem of exact marginalization is that it introduces new edges between all pairs of variables that are directly related to the eliminated variables, adding a so-called elimination clique to the graph, see for example [6]. This, unfortunately, destroys the natural sparsity pattern that is typical to SLAM problems.

Therefore, we apply an approximate marginalization scheme to maintain sparsity, which is important for long-term mapping tasks. The key idea is to replace the elimination clique, which is created when marginalizing a node, by a tree-shaped approximation of the clique. The optimal tree-shaped approximation with respect to the Kullback-Leibler divergence is given by the Chow-Liu tree. Here, the Chow-Liu tree is the maximum-weight spanning tree of the mutual information graph of the clique. This tree can be computed by assigning the mutual information of each two variables to the corresponding edges belonging to the elimination clique and then applying Kruskal’s algorithm. The mutual information is computed according to Davison [3], which describes an efficient solution for the Gaussian case. We refer the reader to [14] for more details on the approximate marginalization scheme.

It should be noted that only the elimination clique is transformed to a tree structure, not the whole pose graph. Otherwise, all loop closing information would be lost.

6 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. In addition to that, we applied our method to a series of popular benchmark datasets. We compare our approach with the performance when no scans are discarded (referred to as “standard approach”). The experiments are designed to show that our approach to informed pose graph compression is well suited for long-term mobile robot mapping as well as for standard SLAM problems.

6.1 Mapping Results

Four different datasets have been considered for the evaluation. We used one self-recorded dataset in which the robot traveled in our lab environment for an extended period of time (Fig. 3), a previously recorded dataset from a Freiburg computer science building (Fig. 4), as well as the Intel Research Lab (Fig. 5), and the FHW dataset (Fig. 1), both provided by Dirk Hähnel.

Fig. 3 shows four maps built during our experiments with a fixed limit to 200 nodes. The first one depicts the pose graph obtained with the standard approach. The second one shows the state of our approach before the robot entered the left side of the corridor. Therefore, 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.

Further mapping results showing the results of our approach in contrast to the standard approach are depicted in Fig. 4 and 5 as well as in the motivating example in Fig. 1. By visual inspection, the obtained grid maps look similar, although only a fraction of the original laser scans have been used to build the grid maps.

To compare the output of the SLAM algorithm more quantitatively, we also analyzed the estimate of the pose graph for the resulting nodes. We especially analyzed the mean and uncertainty estimates for the individual poses of the robot and compared them to the corresponding ones built without compressing the pose graph. Fig. 6 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.



Fig. 3: The robot moved around in an office environment for an extended period of time, visiting the rooms and the corridor many times. 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 for the mutual information, 148 laser scans, 250 edges.



Fig. 4: Obtained map and pose graph for the FR101 dataset (top: standard approach, bottom: our approach). Since the robot does not frequently re-traverse known terrain, few scans were discarded (200 vs. 408 nodes and 246 vs. 723 edges).




Fig. 5: Intel Research Lab. Left: Standard approach, 1802 laser scans, 3916 edges. Right: Our approach preserves the sparsity of the pose graph, 250 laser scans, 349 edges. Arrows indicate locally blurred areas or small alignment errors in the map obtained by the standard approach. In contrast to that, exact marginalization would result in 250 laser scans as well, but in 13052 edges.

6.2 Memory and Runtime Requirements

In this section, we analyze the memory requirements of our approach in terms of the size of the resulting pose graph. In the first experiment, the robot moved around in our lab environment for an extended period of time (Fig. 3). The plots in Fig. 7 and 8 clearly suggest that the experiment leads to an explosion in terms of memory requirements when using the standard approach. This has a direct influence on the computational complexity: First, the optimization technique scales with the number




Fig. 6: Intel Research Lab.: 3σ covariance ellipses of the poses of our approach (red) and the standard one (blue). Our estimates are typically larger since less observations are used.

of edges, which grows roughly quadratically since the robot moves in the same environment and is not exploring new terrain. 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. In contrast to the standard approach, our approach compresses the pose graph such that the number of nodes in the graph remains constant (in Fig. 7 the threshold was set to 200 nodes). If we set a threshold for the mutual information instead of an upper bound for the number of nodes, the complexity does not grow as long as the robot does not explore new territory (see Fig. 8).

Our approach saves computational time as mentioned above but also introduces an overhead through the information-theoretic node selection. This overhead, however, is typically bounded since the number of nodes that have to be considered is bounded since our algorithm runs online and constantly discards nodes. Our current implementation of the information-theoretic laser scan selection is not optimized for speed—significant improvements could be obtained by caching results. Depending on the chosen parameters (particularly l , see Section 4.2) and on the environment that is mapped, the speed of our compression 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.




Fig. 7: Results of an experiment in which the robot moved around in an office environment for an extended period of time. The total number of nodes was restricted to 200.




Fig. 8: Results of an experiment in which the robot moved around in an office environment for an extended period of time. In this experiment, our algorithm discarded all laser scans whose expected information gain was below a threshold.

6.3 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.

6.4 Scan Alignment and Map Quality

We furthermore evaluated the effects of our pose graph compression technique that is applied during mapping on the quality of the resulting grid maps. In theory, the

more observations are available, the better is the estimate. Ignoring measurements will lead to a belief with higher uncertainty. However, today's occupancy grid-based mapping systems typically involve some form of scan alignment or scan matching (e.g. to extract constraints). Such systems have the following disadvantage when it comes to long-term map learning. Whenever the robot obtains a 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 a small alignment error is nonzero. A scan which 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. 5 and 3 depict the maps and graphs obtained from the Intel Research Lab dataset and the long-term experiments conducted in our office environment. The 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 method discards scans and thus produces maps with sharp obstacle boundaries even in cases in which the robot frequently re-traverses already visited places. Although we do not claim that such a sharp map is a better estimate of the world, it better supports the scan matcher and reduces the risk of divergence in the mapping process.

7 Conclusion

In this paper, we presented a method for information-theoretic compression of pose graphs in graph-based SLAM, which is an important step towards long-term mapping. 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 amount of information that a laser scan is expected to provide. Our approach estimates the expected information gain of laser measurements with respect to the resulting occupancy grid map. Real world experiments illustrate the effectiveness of our method for computing compressed pose graphs in the context of graph-based SLAM.

Acknowledgement

This work has partly been supported by the DFG under SFB/TR-8 and by the EC under FP7-231888-EUROPA. We would like to thank Giorgio Grisetti for his contribution to the marginalization as well as Maximilian Beinhofer for fruitful

discussions. Furthermore, thanks to Dirk Hähnel for providing the FHW and the Intel Research Lab datasets.

References

1. M. Bosse, P. M. Newman, J. J. Leonard, and S. Teller. An ATLAS framework for scalable mapping. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, pages 1899–1906, 2003.
2. M. Cummins and P. Newman. FAB-MAP: Probabilistic localization and mapping in the space of appearance. *Int. Journal of Robotics Research*, 27(6):647–665, 2008.
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. C. Estrada, J. Neira, and J.D. Tardós. Hierarchical SLAM: Real-time accurate mapping of large environments. *IEEE Transactions on Robotics*, 21(4):588–596, 2005.
6. 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.
7. J. Folkesson, P. Jensfelt, and H. Christensen. Vision SLAM in the measurement subspace. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, pages 325–330, 2005.
8. 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.
9. V. Ila, J.M. Porta, and J. Andrade-Cetto. Information-based compact pose slam. *IEEE Transactions on Robotics*, 26(1):78–93, 2010.
10. 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.
11. K. Konolige and M. Agrawal. FrameSLAM: From bundle adjustment to realtime visual mapping. *IEEE Transactions on Robotics*, 24(5):1066–1077, 2008.
12. 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.
13. A. Krause and C. Guestrin. Near-optimal nonmyopic value of information in graphical models. In *Proc. of Uncertainty in Artificial Intelligence (UAI)*, 2005.
14. H. Kretzschmar, C. Stachniss, and G. Grisetti. Efficient information-theoretic graph pruning for graph-based SLAM with laser range finders. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, San Francisco, CA, USA, 2011. Accepted for publication.
15. F. Lu and E. Milios. Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, 4:333–349, 1997.
16. D.J.C. MacKay. *Information theory, inference, and learning algorithms*. Cambridge Univ Press, 2003.
17. K. Ni, D. Steedly, and F. Dellaert. Tectonic SAM: Exact; out-of-core; submap-based slam. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, Rome, Italy, 2007.
18. E. Olson. *Robust and Efficient Robotic Mapping*. PhD thesis, MIT, Cambridge, MA, USA, June 2008.
19. P. Pfaff, R. Triebel, C. Stachniss, P. Lamon, W. Burgard, and R. Siegwart. Towards mapping of cities. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, Rome, Italy, 2007.
20. 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.