Auton Robot (2006) 21:103–122
DOI 10.1007/s10514-006-9043-2

Treemap: An $O(\log n)$ algorithm for indoor simultaneous localization and mapping

Udo Frese

Accepted: 27 January 2006 / Published online: 25 August 2006
© Springer Science + Business Media, LLC 2006

Abstract This article presents a very efficient SLAM algorithm that works by hierarchically dividing a map into local regions and subregions. At each level of the hierarchy each region stores a matrix representing some of the landmarks contained in this region. To keep those matrices small, only those landmarks are represented that are observable from outside the region.

A measurement is integrated into a local subregion using $O(k^2)$ computation time for k landmarks in a subregion. When the robot moves to a different subregion a full least-square estimate for that region is computed in only $O(k^3 \log n)$ computation time for n landmarks. A global least square estimate needs $O(kn)$ computation time with a very small constant (12.37 ms for $n = 11300$).

The algorithm is evaluated for map quality, storage space and computation time using simulated and real experiments in an office environment.

Keywords Mobile robots · SLAM · Information matrix · Hierarchical decomposition

1. Introduction

The problem of building a map from local observations of the environment is a very old one, as old as maps themselves. While geodesy, the science of surveying in general, dates back to 8000 B.C., it was C.F. Gauss who first formalized the problem from the perspective of statistical estimation in his article "*Theoria combinationis observationum erroribus minimis obnoxiae*"[1] (1821).

In the much younger realm of robotics, the corresponding problem is that of simultaneous localization and mapping (SLAM). It requires the robot to continuously build a map from sensor data while exploring the environment. It has been a subject of research since the mid 1980s gaining enormous popularity in recent years. Most approaches adhere to the Gaussian formalization. They estimate a vector of n features, e.g. landmarks or laser scan reference frames, by minimizing a quadratic error function, i.e. by implicitly solving a linear equation system. With this well established methodology the main question is how to compute or approximate the estimate efficiently. To make this more explicit, we have proposed three important requirements which an ideal SLAM algorithm should fulfill (Frese and Hirzinger, 2001; Frese, 2006a).

(R1) Bounded uncertainty. *The uncertainty of any aspect of the map should not be much larger than the minimal uncertainty that could theoretically be derived from the measurements.*

(R2) Linear storage space. *The storage space of a map covering a large area should be linear in the number of landmarks.*

(R3) Linear update cost. *Incorporating a measurement into a map covering a large area should have a computational cost at most linear in the number of landmarks.*

This article is based on the authors studies at the German Aerospace Center.

U. Frese
FB 3, Mathematik und Informatik, SFB/TR 8 Spatial Cognition, Universität Bremen,
Bremen, Germany
e-mail: ufrese@informatik.uni-bremen.de

[1] "Theory of the combination of observations least subject to error."

(R1) binds the map to reality and limits approximations. (R2) and (R3) regard efficiency, requiring linear space and time consumption. We feel that one should aim at (R1) rather than at the much weaker criterion of asymptotic convergence. Even after going through the environment only a single time, a useful map is desirable. Most sensor data allows this and, according to (R1) so should the SLAM algorithm.

The contribution of this article is *treemap*, a hierarchical SLAM algorithm that meets the requirements (R1)–(R3). It works by dividing the map into regions and subregions. When integrating a measurement it needs $O(k^2)$ computation time for updating the estimate for a region with k landmarks, $O(k^3 \log n)$ when the robot moves to a different region, and $O(kn)$ to compute an incremental estimate for the whole map. The algorithm is landmark based and requires known data association. It has two drawbacks. First, it requires a "topologically suitable building" (Section 4); and second, its implementation is relatively complex.

The article is organized as follows. After a brief review of related work (Section 2), we derive the algorithm (Sections 3–6). It follows with a comparison to the closely related Thin Junction Tree Filter (Section 7) by Paskin (2003) and an investigation of map quality and computation time based on simulations (Section 8) and experiments in a 60 × 45 m office building (Section 9). A companion technical report (Frese, 2006b) supplements the discussion with an improved—but more complicated—method for passing the robot pose between regions, a nonlinear extension, and the algorithm's pseudocode.

2. State of the art

After the fundamental article by Smith et al. in 1988 most work on SLAM was based on the Extended Kalman Filter (EKF) that allows SLAM to be treated as an estimation problem in a theoretical framework. However, the problem of large computation time remained. The EKF maintains the posterior distribution of the robot pose and n landmarks as a $3+2n$ dimensional Gaussian with correlations between all landmarks. This is essentially for SLAM but requires to update the EKF's covariance matrix after each measurement, taking $O(n^2)$ time. This limited the use to the order of hundreds of landmarks.

Recently, interest in SLAM has increased dramatically and several more efficient algorithms have been developed. Many approaches exploit the fact that observations are *local* in the sense that from a single robot pose only a few (k) landmarks are visible. In the following the more recent contributions will be briefly reviewed (Table 1). An overview is given by Thrun et al. (2005) and a discussion by Frese (2006a).

To meet requirement (R1) an algorithm must maintain some form of correlations in the whole map. To my knowledge the first SLAM algorithm achieving this with a computation time below $O(n^2)$ per measurement was the relaxation algorithm (Duckett et al., 2000, 2002). The algorithm employs an iterative equation solver called *relaxation* to the linear equation system appearing in maximum likelihood estimation. One iteration is applied after each measurement with $O(kn)$ computation time and $O(kn)$ storage space. After closing a loop, more iterations are necessary leading to $O(kn^2)$ computation time in the worst case. This was later improved by the Multilevel Relaxation (MLR) algorithm (Frese et al., 2004). This algorithm optimizes the map at different levels of resolution, leading to $O(kn)$ computation time.

Montemerlo et al. (2002) derived an algorithm called *FastSLAM* from the observation that the landmark estimates are conditionally independent, given the robot pose. Basically, the algorithm is a particle filter (M particles) in which every particle represents a sampled robot trajectory plus a set of n Kalman filters estimating the landmarks conditioned on the trajectory. The number of particles M is a difficult trade-off between computation time and map quality. However,

Table 1 Performance of different SLAM algorithms with n landmarks, m measurements, p robot poses and k landmarks local to the robot (cf. Section 2). UDA stands for 'Uncertain Data Association'. A √ means the algorithm can handle landmarks with uncertain identity. A C means covariance is available for performing χ^2 tests

	(R1)			(R2)	(R3)		
	UDA	Non-linear	Map quality	Memory	Update	Global update	Loop
ML	C	√	√	m	....	$(n+p)^3$	....
EKF	C		√	n^2	....	n^2	....
CEKF	C		√	$n^{\frac{3}{2}}$	k^2	.... $kn^{\frac{3}{2}}$	....
Relaxation		√	√	kn	....	kn	kn^2
MLR		√	√	kn	....	kn	....
FastSLAM	√	√	see Section 2	Mn	....	$M \log n$	....
SEIF				kn	....	k^2	....
w. full update			√	kn	....	kn	kn^2
TJTF	C	√	√	$k^2 n$	k^3	.... $k^3 n$	....
Treemap	C	√	√	kn	k^2	...... $k^3 \log n$	......
w. global map	C	√	√	kn	....	kn	....

the algorithm can handle uncertain landmark identification (Nieto et al., 2003), which is a unique advantage over the other algorithms discussed. Later Eliazar and Parr (2003) as well as Stachniss and Burgard (2004) extended the framework to using plain evidence grids as particles (in a compressed representation). Their approach constructs maps in difficult situations without landmark extraction or scan matching.

Guivant and Nebot (2001, 2003) developed the *Compressed EKF* (*CEKF*) that allows the accumulation of measurements in a local region with k landmarks at cost $O(k^2)$ independent from n. When the robot leaves this region, the accumulated result is propagated to the full EKF (*global update*) at cost $O(kn^2)$. The global update can be approximated more efficiently in $O(kn^{3/2})$ with $O(n^{3/2})$ storage space needed.

Thrun et al. (2004) presented a "constant time" algorithm called the *Sparse Extended Information Filter (SEIF)*, which represents uncertainty with an information matrix instead of a covariance matrix. The algorithm exploits the observation that the information matrix is approximately sparse requiring $O(kn)$ storage space. This property has later been proven by Frese (2005). To compute a map estimate, a system of n linear equations has to be solved. Thrun et al. use relaxation but update only $O(k)$ landmarks after each measurement (using so-called amortization). This can derogate map quality, since in the numerical literature, relaxation is reputed to need $O(kn^2)$ time for reducing the equation error by a constant factor (Press et al., 1992). Imagine closing a loop of length n and going around that loop a second time. Still the estimate will not have reasonably converged because only $O(k^2n)$ time has been spent. If all landmarks are updated each step *(SEIF w. full update)* performance is asymptotically the same as with relaxation.

Leonard and Feder (2001) avoid the problem of updating an estimate for n landmarks by dividing the map into submaps. Their *Decoupled Stochastic Mapping (DSM)* approach represents each submap in global coordinates by an EKF and updates only the current local submap. The approach is very fast ($O(k^2)$) but, as they note, can introduce overconfidence when passing the robot pose between submaps.

Bosse et al. (2004) in contrast make submaps probabilistically independent in their *Atlas* framework by using a local reference frame. Then, links between adjacent frames are derived by matching local maps. From the resulting graph an estimate is computed. A similar approach is taken by Estrada et al. (2005). Both systems are heterogenous. On the local level a full least square solution is obtained. But on the global level the complex probabilistic relation between submaps is aggregated into a single 3-DOF link between their reference frames.

Paskin (2003) derived the *Thin Junction Tree Filter* (TJTF) from viewing the problem as a Gaussian graphical model. This approach is closely related to treemap although both have been independently derived from different perspectives. They are compared in Section 7.

In the following the treemap algorithm will be introduced. It can be used in the same way as CEKF providing an estimate for k landmarks of a local region but with only $O(k^3 \log n)$ computation time when changing the region instead of $O(kn^{3/2})$ for CEKF. Alternatively the algorithm can also compute a global estimate for all n landmarks. Computation time is then $O(kn)$, but with a constant so small that this can be done for almost "arbitrarily" large maps (12.37 ms for $n = 11300$, Intel Xeon 2.7 GHz). This is the main contribution from a practical perspective. Note, that while treemap can also provide covariance information, this requires additional computation in contrast to CEKF.

3. Treemap data structure

3.1. Motivating idea

We will first discuss the general idea that motivates the treemap approach. The description differs slightly from the actual implementation but provides a large-picture understanding of the algorithm.

Imagine the robot is in a building (Fig. 1a) that is virtually divided into two parts A and B. Now consider the following question here:

If the robot is in part A, *what is the information needed about* B?

Some of B's landmarks are observable from A and involved in measurements while the robot is in A. The algorithm must have all previously gathered information about these landmarks explicitly available. This information is more than just the measurements that directly involve those landmarks. Rather all measurements in B can indirectly contribute to the information. So probabilistically speaking, the information needed about B is the marginal distribution of landmarks observable both from A and from B conditioned on measurements taken in B.[2]

The idea can be applied recursively by dividing the building into a binary[3] tree of regions (Fig. 1). The recursion stops when the size of a region is comparable to the robot's field of view.

The marginal distribution for a region can be computed recursively. The marginals for the two subregions are multiplied and landmarks are marginalized out that are

[2] This only holds strictly when there is no odometry (cf. Section 5).

[3] Using a binary hierarchy simplifies bookkeeping.

(a)

(b)

Fig. 1 *Geometric view*. (a) A building hierarchically decomposed in two levels (L1, L2) and (b) respective tree representation with 7 nodes ($\mathbf{n}_{1...7}$). The region corresponding to a node is shown next to the node. Landmarks inside a region that are visible from outside are listed in the node. Only the marginal distribution of these landmarks is needed if the robot is outside the respective region

not visible from the outside of that larger region anymore. This core computation is the same as employed by TJTF. The key benefit of this approach is that for integrating a measurement only the region containing the robot and its super-regions need to be updated. All other regions remain unaffected.

Treemap consists of three parts: Core propagation of information in the tree (Sections 3–4); preprocessing to get information into the tree (Section 5); and hierarchical tree partitioning to find a good tree (Section 6).

3.2. Formal Bayesian view

In a preprocessing step the original measurements are converted into probabilistic constraints $p(X \mid z_i)$ on the state vector of landmark positions X. At the moment, let us take an abstract probabilistic perspective as to how treemap computes an estimate $\hat{x}$ from these constraints. We will subsequently describe the Gaussian implementation as well as how to get the constraints z_i from the original measurements.

The constraints are assigned to leaves of the tree with the intention to group constraints that share landmarks. With respect to the motivating idea each leaf defines a local region and correspondingly each inner node a super-region. However formally a node $\mathbf{n}$ just represents the set of constraints assigned to leaves below $\mathbf{n}$ without any explicit geometric definition. For a node $\mathbf{n}$, the left and right child and the parent are denoted by $\mathbf{n}_{\swarrow}$, $\mathbf{n}_{\searrow}$ and $\mathbf{n}_{\uparrow}$, respectively. We often have to deal with subsets of observations or landmarks according to where they are represented within the tree relative to the node $\mathbf{n}$ (Fig. 2). Thus, let $z[\mathbf{n}{:}\downarrow]$, $z[\mathbf{n}{:}\swarrow]$, $z[\mathbf{n}{:}\searrow]$, and $z[\mathbf{n}{:}\uparrow]$ denote the constraints assigned to leaves *below* ($\downarrow$), *left-below* ($\swarrow$), *right-below* ($\searrow$), and *above* ($\uparrow$) node $\mathbf{n}$, respectively. The term *above* $\mathbf{n}$ refers to *all regions outside the subtree below* $\mathbf{n}$(!). As a special case, for a leaf $\mathbf{n}$, let $z[\mathbf{n}{:}\swarrow\searrow\updownarrow]$ denote all constraints at $\mathbf{n}$. Analogous expressions $X[\mathbf{n}{:}\ldots]$ denote the landmarks involved in the corresponding constraints $z[\mathbf{n}{:}\ldots]$. While constraint sets for different directions $\{\swarrow, \searrow, \uparrow\}$ are disjoint, the corresponding landmark sets may overlap because different constraints may involve the same landmark. These shared features dictate the computations at node $\mathbf{n}$ as the tree is updated. With the assumptions presented in Section 4, their number is small ($O(k)$) and, thus, the overall computation involves many low-dimensional distributions instead of one high-dimensional.

3.3. Update (upwards)

Figure 3 depicts the data flow in treemap that consists of integration ($\odot$) and marginalization (Ⓜ), i.e. multiplying and factorizing probability distributions. We will now derive and prove the computation.

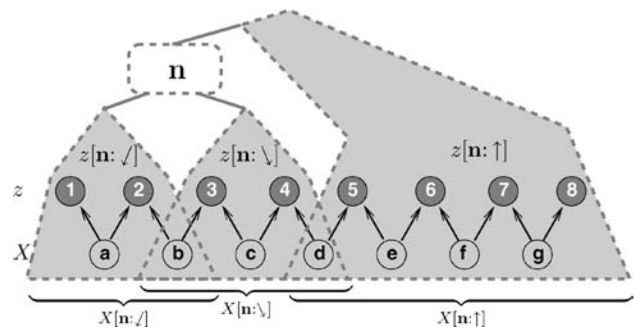

Fig. 2 *Bayesian view*. In this example landmarks $x_{\mathsf{a}\ldots\mathsf{g}}$ are observed. They are connected by constraints $z_{2\ldots7}$ between consecutive landmarks and two absolute constraints z_1, z_8. The arrows and circles show this probabilistic input as a Bayes net with observed nodes in gray. The dashed outlines illustrate the information from the view of a single node $\mathbf{n}$. It divides the tree into three parts, *left-below* $\swarrow$, *right-below* $\searrow$ and *above* $\uparrow$ (more precisely not below). Hence the constraints z are disjointly divided into $z[\mathbf{n}{:}\swarrow] = z_{1\ldots2}$, $z[\mathbf{n}{:}\searrow] = z_{3\ldots4}$ and $z[\mathbf{n}{:}\uparrow] = z_{5\ldots8}$. The corresponding landmarks $X[\mathbf{n}{:}\swarrow] = X_{\mathsf{a}\ldots\mathsf{b}}$, $X[\mathbf{n}{:}\searrow] = X_{\mathsf{b}\ldots\mathsf{d}}$ and $X[\mathbf{n}{:}\uparrow] = X_{\mathsf{d}\ldots\mathsf{g}}$ however overlap ($X[\mathbf{n}{:}\swarrow\searrow] = X_{\mathsf{b}}$, $X[\mathbf{n}{:}\uparrow\searrow] = x_{\mathsf{d}}$). The key insight is, that $X[\mathbf{n}{:}\downarrow\uparrow] = X[\mathbf{n}{:}\swarrow\uparrow \vee \searrow\uparrow] = X_{\mathsf{d}}$ separates the constraints $z[\mathbf{n}{:}\downarrow]$ and landmarks $X[\mathbf{n}{:}\downarrow\updownarrow]$ below $\mathbf{n}$ from the constraints $z[\mathbf{n}{:}\uparrow]$ and landmarks $X[\mathbf{n}{:}\updownarrow\uparrow]$ above $\mathbf{n}$, so both are conditionally independent given $X[\mathbf{n}{:}\downarrow\uparrow]$.

As input, treemap receives a distribution $p^I_{\mathbf{n}}$ defined as $p(X[\mathbf{n}{:}\downarrow]|z[\mathbf{n}{:}\downarrow])$ at each leaf. It is computed from the probabilistic model for the constraints assigned to $\mathbf{n}$. The output is the integrated information $p_{\mathbf{n}} = p(X[\mathbf{n}{:}\downarrow]|z)$ at each leaf. During the computation, intermediate distributions $p^M_{\mathbf{n}}$ and $p^C_{\mathbf{n}}$ are passed through the tree and stored at the nodes, respectively. In general, $p^I_{\mathbf{n}}$, $p^M_{\mathbf{n}}$, $p^C_{\mathbf{n}}$, and $p_{\mathbf{n}}$ refer to distributions actually computed by the algorithm, whereas all distributions $p(X[\ldots]|z[\ldots])$ refer to *the* distribution of the landmarks $X[\ldots]$ given the constraints $z[\ldots]$ according to the abstract probabilistic input model shown in Fig. 2. With this notion proving the algorithm means to derive equations $p^{\cdots}_{\mathbf{n}} = p(X[\ldots]|z[\ldots])$ expressing that the computed result equals the desired distribution from the input model.

Let us first consider the update operation that begins at the leaves and is recursively applied upwards. The update computes the marginal $p^M_{\mathbf{n}}$ and conditional $p^C_{\mathbf{n}}$ either from the input distribution $p^I_{\mathbf{n}}$ or from the children's marginals $p^M_{\mathbf{n}_{\swarrow}}$ and $p^M_{\mathbf{n}_{\searrow}}$.

$$p^M_{\mathbf{n}} = p(X[\mathbf{n}{:}\ \downarrow\uparrow]|z[\mathbf{n}{:}\downarrow]) \quad (1)$$

$$p^C_{\mathbf{n}} = p(X[\mathbf{n}{:}\swarrow\searrow\updownarrow]|X[\mathbf{n}{:}\downarrow\uparrow], z). \quad (2)$$

The marginal distribution (1) describes the posterior for the landmarks both above and below $X[\mathbf{n}{:}\downarrow\uparrow]$ conditioned upon the constraints $z[\mathbf{n}{:}\downarrow]$ below $\mathbf{n}$. These landmarks are by definition also involved in constraints which are not yet

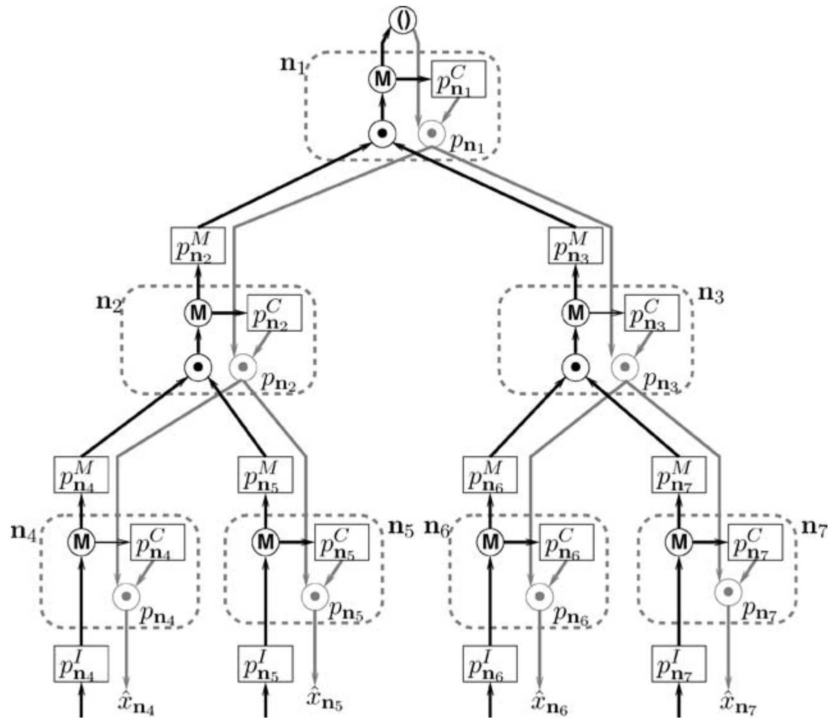

Fig. 3 *Data flow view*. The probabilistic computations performed in the tree shown in Fig. 1(b). The leaves store the input constraints $p^I_{\mathbf{n}}$. During updates (black arrows) a node **n** integrates ($\odot$) the distributions $p^M_{\mathbf{n}_\swarrow}$ and $p^M_{\mathbf{n}_\searrow}$ passed by its children. Then the result is factorized (Ⓜ) as the product of a marginal p^M_n passed to the parent and a conditional p^C_n stored at the node. To compute an estimate (gray arrows) each node **n** receives a distribution $p_{\mathbf{n}_\uparrow}$ from its parent, integrates ($\odot$) it with the conditional $p^C_{\mathbf{n}}$, and passes the result $p_{\mathbf{n}}$ down to its children. In the end estimates $\hat{x}_{\mathbf{n}}$ for all landmarks are available at the leaves

integrated into $p^M_{\mathbf{n}}$. So $p^M_{\mathbf{n}}$ is passed to the parent for further processing. In contrast, $p^C_{\mathbf{n}}$ contains those landmarks $X[\mathbf{n}{:}\swarrow\searrow\not\uparrow]$ for which **n** is the least common ancestor of all constraints involving them. These constraints have already been integrated, so $p^C_{\mathbf{n}}$ needs no more processing and can be finally stored at **n**. Overall, a landmark is passed upwards in $p^M_{\mathbf{n}}$ up to the node where all constraints involving that landmark have been integrated and then it is stored in $p^C_{\mathbf{n}}$.

We now derive the recursive computation of $p^M_{\mathbf{n}}$ and $p^C_{\mathbf{n}}$ proving (1) and (2) by induction. An inner node **n** multiplies ($\odot$) the marginals $p^M_{\mathbf{n}_\swarrow}$ and $p^M_{\mathbf{n}_\searrow}$ passed by its children. Assuming (1) for $\mathbf{n}_\searrow$ we get

$$p^M_{\mathbf{n}_\searrow} = p(X[\mathbf{n}_\searrow : \downarrow\uparrow]|z[\mathbf{n}_\searrow : \downarrow]). \tag{3}$$

Being above $\mathbf{n}_\searrow$ means either to be above **n** or to be left-below **n**.

$$= p(X[\mathbf{n}{:}\searrow\uparrow \vee \swarrow\searrow\not\uparrow]|z[\mathbf{n}{:}\searrow]) \tag{4}$$

To multiply $p^M_{\mathbf{n}_\swarrow}$ and $p^M_{\mathbf{n}_\searrow}$ we must formally interpret both as a distribution for the union of landmarks. This is possible, since $X[\mathbf{n}{:}\swarrow\not\searrow\uparrow]$ are by definition not involved in $z[\mathbf{n}{:}\searrow]$ at all.

$$= p(X[\mathbf{n}{:}\searrow\uparrow \vee \swarrow\searrow\not\uparrow \vee \swarrow\not\searrow\uparrow]|z[\mathbf{n}{:}\searrow]) \tag{5}$$

$$= p\big(X[\mathbf{n}{:}\downarrow\uparrow \vee \swarrow\searrow\not\uparrow]|z[\mathbf{n}{:}\searrow]) \tag{6}$$

This argument may appear rather technical at first sight but it ensures that in defining $p^M_{\mathbf{n}}$ by (1) we actually found that part of $p(X|z[\mathbf{n}{:}\downarrow])$ that cannot be fully processed below **n** and has to be passed to the parent. Certainly a symmetric result holds for $p^M_{\mathbf{n}_\swarrow}$, so both can be multiplied ($\odot$) gathering all information below **n**.

$$p^M_{\mathbf{n}_\swarrow} \cdot p^M_{\mathbf{n}_\searrow} = p(X[\mathbf{n}{:}\downarrow\uparrow \vee \swarrow\searrow\not\uparrow]|z[\mathbf{n}\swarrow])$$

$$\cdot p(X[\mathbf{n}{:}\downarrow\uparrow \vee \swarrow\searrow\not\uparrow]|z[\mathbf{n}{:}\searrow]) \tag{7}$$

$$= p(X[\mathbf{n}{:}\downarrow\uparrow \vee \swarrow\searrow\not\uparrow]|z[\mathbf{n}\downarrow]) \tag{8}$$

Let $Y = X[\mathbf{n}{:}\downarrow\uparrow \vee \swarrow\searrow\not\uparrow]$ be the vector of landmarks involved in $p^M_{\mathbf{n}_\swarrow}$ or $p^M_{\mathbf{n}_\searrow}$. Treemap divides $Y = \binom{U}{V}$ into those landmarks $V = X[\mathbf{n}{:}\downarrow\uparrow]$ involved in constraints above **n** and those $U = X[\mathbf{n}{:}\swarrow\searrow\not\uparrow]$ for which **n** is the least common ancestor of all constraints involving them. Landmarks U are marginalized out (Ⓜ) factorizing the distribution as the product

$$p^M_{\mathbf{n}_\swarrow}\begin{pmatrix}u\\v\end{pmatrix} \cdot p^M_{\mathbf{n}_\searrow}\begin{pmatrix}u\\v\end{pmatrix} = p^M_{\mathbf{n}}(v) \cdot p^C_{\mathbf{n}}(u \mid v) \tag{9}$$

of the marginal $p^M_{\mathbf{n}}(V)$ passed to the parent and the conditional $p^C_{\mathbf{n}}(U|V)$ stored at **n**. We can verify, that the computed $p^M_{\mathbf{n}}$ satisfies (1) and $p^C_{\mathbf{n}}$ satisfies (2) by marginalizing respectively conditioning both sides of (8).

$$p^M_{\mathbf{n}} = p(X[\mathbf{n}{:}\downarrow\uparrow]|z[\mathbf{n}\downarrow]) \tag{10}$$

$$p^C_{\mathbf{n}} = p(X[\mathbf{n}{:}\swarrow\searrow\not\uparrow] | X[\mathbf{n}{:}\downarrow\uparrow], z[\mathbf{n}\downarrow]) \tag{11}$$

$$= p(X[\mathbf{n}{:}\swarrow\searrow\not\uparrow] | X[\mathbf{n}{:}\downarrow\uparrow], z) \tag{12}$$

The second step (12) is the formal key point of the overall approach. In Bayes net terminology (Fig. 2) $X[\mathbf{n}{:}\downarrow\uparrow]$ separates the constraints $z[\mathbf{n}{:}\downarrow]$ and landmarks $X[\mathbf{n}{:}\downarrow\not\uparrow]$ below $\mathbf{n}$ from the constraints $z[\mathbf{n}{:}\uparrow]$ and landmarks $X[\mathbf{n}{:}\not\downarrow\uparrow]$ above $\mathbf{n}$. So $X[\mathbf{n}{:}\swarrow\searrow\not\uparrow]$, which is part of $X[\mathbf{n}{:}\downarrow\not\uparrow]$, is conditionally independent from the remaining constraints $z[\mathbf{n}{:}\uparrow]$.

Now we have established that if (1) holds for a given nodes children, then (1) and (2) hold for $p^M_{\mathbf{n}}$ and $p^C_{\mathbf{n}}$ computed by the node. We still have to verify these equations for leaves. Then by induction they hold for all nodes. At a leaf $\mathbf{n}$ all original constraints that were assigned to that leaf are multiplied and stored as input distribution

$$p^I_{\mathbf{n}} = \prod_{i \text{ assigned to } \mathbf{n}} p(X|z_i) = p(X[\mathbf{n}{:}\downarrow]|z[\mathbf{n}{:}\downarrow]) \tag{13}$$

$$= p(X[\mathbf{n}{:}\downarrow\uparrow \vee \swarrow\searrow\not\uparrow]|z[\mathbf{n}{:}\downarrow]). \tag{14}$$

For a leaf $\mathbf{n}$ we defined $X[\mathbf{n}{:}\swarrow\searrow]$ as those landmarks involved in constraints assigned to that leaf. So for a leaf, $p^I_{\mathbf{n}}$ satisfies the same condition (8) as $p^M_{\mathbf{n}\swarrow} \cdot p^M_{\mathbf{n}\searrow}$ for an inner node. Hence, after marginalization (1) and (2) hold for leaves with the same arguments as for inner nodes.

As a final remark, $p^M_{\mathbf{root}} = ()$ is empty by (1), because there is nothing above **root**. So it is the end of the upward update-arrows and the start of the downward state-recovery arrows (Fig. 3).

3.4. State recovery (downwards)

Now let us consider how to compute a state estimate from the $p^C_{\mathbf{n}}$ (gray arrows pointing downwards). Here the goal is that every node $\mathbf{n}$ passes

$$p_{\mathbf{n}} = p(X[\mathbf{n}{:}\downarrow\uparrow \vee \swarrow\searrow\not\uparrow]|z) \tag{15}$$

down. Hence a leaf computes the marginal of landmarks involved since $X[\mathbf{n}{:}\downarrow\uparrow \vee \swarrow\searrow\not\uparrow]$ equals $X[\mathbf{n}{:}\downarrow]$. The final estimate $\hat{x}_{\mathbf{n}}$ is computed as

$$\hat{x}_{\mathbf{n}} = E(p_{\mathbf{n}}) = E(X[\mathbf{n}{:}\downarrow]|z). \tag{16}$$

Since every update changes $p_{\mathbf{n}}$, it is computed on the fly and not stored.

Now we derive (15) by induction. Let us assume a node $\mathbf{n}$ receives

$$p_{\mathbf{n}_\uparrow} = p(X[\mathbf{n}_\uparrow{:}\downarrow\uparrow \vee \swarrow\searrow\not\uparrow]|z) \tag{17}$$

from its parent. A landmark below $\mathbf{n}_\uparrow$ is either below $\mathbf{n}$ or below the sibling of $\mathbf{n}$. The latter ones are marginalized out resulting in

$$p(X[\mathbf{n}{:}\downarrow\uparrow]|z). \tag{18}$$

This step is not shown in Fig. 3 because it is implicitly done in the actual Gaussian implementation (cf. Section 3.5). The result is multiplied ($\odot$) with the conditional $p^C_{\mathbf{n}}$ stored at $\mathbf{n}$ and passed downwards as $p_{\mathbf{n}}$.

$$p_{\mathbf{n}} = p(X[\mathbf{n}{:}\downarrow\uparrow]|z) \cdot p^C_{\mathbf{n}} \tag{19}$$

$$= p(X[\mathbf{n}{:}\downarrow\uparrow]|z) \cdot p(X[\mathbf{n}{:}\swarrow\searrow\not\uparrow]|X[\mathbf{n}{:}\downarrow\uparrow], z) \tag{20}$$

$$= p(X[\mathbf{n}{:}\downarrow\uparrow \vee \swarrow\searrow\not\uparrow]|z) \tag{21}$$

We have shown, that if node $\mathbf{n}$ receives $p_{\mathbf{n}\uparrow}$ with (15) it passes a distribution $p_{\mathbf{n}}$ to its children holding (15) too. As the induction start $X[\mathbf{root}{:}\downarrow\uparrow]$ in (18) is empty, so by induction (15) holds for all $p_{\mathbf{n}}$.

3.5. Gaussian implementation

Treemap uses Gaussians for all probability distributions. Thereby the probabilistic computations reduce to matrix operations and the algorithm becomes an efficient linear equation solver for a specific class of equations. The performance is much improved by using different representations for updates ($p^I_{\mathbf{n}}$, $p^M_{\mathbf{n}}$), for state recovery ($p_{\mathbf{n}}$), and for the conditional $p^C_{\mathbf{n}}$ linking both. We will now derive formulas for the three operations $\odot$ (update), Ⓜ, and $\odot$ (state recovery) involved.

Distributions $p^I_{\mathbf{n}}$ and $p^M_{\mathbf{n}}$ are stored in information form as

$$-\log\ p^M_{\mathbf{n}}(y) = y^T A_{\mathbf{n}} y + y^T b_{\mathbf{n}} + \text{const}. \tag{22}$$

3.5.1. Update (upwards)

If treemap is used directly with landmark–landmark constraints, $p^I_{\mathbf{n}}$ is computed as usual by linearizing the constraints, expressing the approximated χ^2 error by an information matrix and vector and adding these for all constraints assigned to the leaf $\mathbf{n}$ (Thrun et al., 2005, Section 11.4.3). In Section 5 we will discuss how to derive $p^I_{\mathbf{n}}$ in a preprocessing step from robot–landmark and robot–robot constraints.

To perform the multiplication $\odot$ at node $\mathbf{n}$, first $(A_{\mathbf{n}\swarrow}, b_{\mathbf{n}\swarrow})$ as well as $(A_{\mathbf{n}\searrow}, b_{\mathbf{n}\searrow})$ are permuted and extended with 0-rows/columns such that the same row/column corresponds to the same landmark in both. Additionally landmarks of $X[\mathbf{n}{:}\swarrow\searrow\not\uparrow]$ are permuted to the upper rows/left columns and

landmarks of $X[\mathbf{n}{:}\downarrow\uparrow]$ to the lower rows/right columns. This will help later for marginalization. Then they are added.

$$-\log\left(p^M_{\mathbf{n}_\swarrow}(y)\ p^M_{\mathbf{n}_\searrow}(y)\right) = -\log p^M_{\mathbf{n}_\swarrow}(y) - \log p^M_{\mathbf{n}_\searrow}(y) \tag{23}$$

$$= y^T(A_{\mathbf{n}_\swarrow} + A_{\mathbf{n}_\searrow})y + y^T(b_{\mathbf{n}_\swarrow} + b_{\mathbf{n}_\searrow}) + \text{const} \tag{24}$$

To perform the marginalization (Ⓜ), $A_{\mathbf{n}_\swarrow} + A_{\mathbf{n}_\searrow}$ is viewed as a 2×2 block matrix and $b_{\mathbf{n}_\swarrow} + b_{\mathbf{n}_\searrow}$ as a 2 block vector

$$= \begin{pmatrix} u \\ v \end{pmatrix}^T \begin{pmatrix} P & R^T \\ R & S \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} + \begin{pmatrix} u \\ v \end{pmatrix}^T \begin{pmatrix} c \\ d \end{pmatrix} + \text{const}. \tag{25}$$

The first block row/column corresponds to landmarks $U = X[\mathbf{n}{:}\swarrow\searrow\not\uparrow]$ to be marginalized out and stored in $p^C_{\mathbf{n}}$. The second block row/column corresponds to landmarks $V = X[\mathbf{n}{:}\downarrow\uparrow]$ to be passed in $p^M_{\mathbf{n}}$. By a straight-forward but rather lengthy calculation it follows that

$$\begin{aligned} &= v^T(S - RP^{-1}R^T)v + v^T(-RP^{-1}c + d) + \text{const} \\ &\quad + (Hv + h - u)^T P(Hv + h - u), \end{aligned} \tag{26}$$

$$\text{with } H = -P^{-1}R^T \text{ and } h = -P^{-1}c/2. \tag{27}$$

The first line of (26) defines a Gaussian for v in information form not involving u at all. The second line defines a Gaussian for u with covariance P^{-1} and mean $Hv + h$. The first does not contribute to the conditional $p(U \mid V)$ and the second not to the marginal $p(v)$. Thus

$$\begin{aligned} -\log p^M_{\mathbf{n}}(v) = v^T(S - RP^{-1}R^T)v + v^T(-RP^{-1}c + d) \\ + \text{const} \end{aligned} \tag{28}$$

$$-\log p^C_{\mathbf{n}}(u|v) = (Hv + h - u)^T P(Hv + h - u) \tag{29}$$

holds. Algorithmically treemap computes the information matrix $A^M_{\mathbf{n}}$ and vector $b^M_{\mathbf{n}}$ of $p^M_{\mathbf{n}}$ by

$$A^M_{\mathbf{n}} = S - RP^{-1}R^T, \quad b^M_{\mathbf{n}} = -RP^{-1}c + d \tag{30}$$

and passes it to the parent node. This is the well known marginalization formula for Gaussians (Thrun et al., 2005, Table 11.6) which is also known as Schur-complement (Horn and Johnson, 1990). Equation (29) is remarkable. It represents $p(u \mid v)$ *in terms of* v as a single Gaussian in u with mean $Hv + h$. For general distributions no such simple relation will hold. Treemap stores $p^C_{\mathbf{n}}$ as (P^{-1}, H, h).

3.5.2. *State recovery (downwards)*

With this representation for $p^C_{\mathbf{n}}$ state recovery ⊙ can be implemented very efficiently in covariance form. The mean $\hat{v} = E(v \mid z)$ is passed by the parent node and the mean of u is correspondingly

$$\hat{y} = \begin{pmatrix} \hat{u} \\ \hat{v} \end{pmatrix},$$

$$\begin{aligned} \hat{u} &= E(u \mid z) \overset{\text{Gaussian}}{=} E(u \mid v = E(v \mid z), z) \\ &= H\hat{v} + h. \end{aligned} \tag{31}$$

Note, that $E(u \mid z) = E(u \mid v = E(v \mid z))$ only holds for Gaussians. In general the full distribution $p(v|z)$ is necessary to compute $E(u|z)$ from $E(u|v, z)$. So for recovering the global state estimate $\hat{x} = E(X|z)$, it suffices to propagate the mean downwards—it is not necessary to propagate covariances at all. In this case, state recovery requires only a single matrix-vector product in each node and is extremely efficient.

If the covariance is desired, it can be propagated the same way. If $\text{cov}(v) = C_v$ is passed by the parent node, $\text{cov}\binom{u}{v}$ can be computed as

$$\begin{aligned} C = \text{cov}\,\hat{y} &= \text{cov}\begin{pmatrix} u \\ v \end{pmatrix} \\ &= \text{cov}\left(\begin{pmatrix} Hv + h \\ v \end{pmatrix} - \begin{pmatrix} Hv + h - u \\ 0 \end{pmatrix}\right) \end{aligned} \tag{32}$$

$$\begin{aligned} &= \begin{pmatrix} HC_vH^T & HC_v \\ C_vH^T & C_v \end{pmatrix} + \begin{pmatrix} P^{-1} & 0 \\ 0 & 0 \end{pmatrix} \\ &= \begin{pmatrix} HC_vH^T + P^{-1} & HC \\ C_vH^T & C_v \end{pmatrix}. \end{aligned} \tag{33}$$

The last equation follows from $p^C_{\mathbf{n}}$ defining a P^{-1} covariance Gaussian on $Hv + h - u$ by (29). This Gaussian is independent from the one passed by the parent node, since the marginalization (Ⓜ) factorizes into two independent distributions $p^M_{\mathbf{n}}$ and $p^C_{\mathbf{n}}$. The result of recursive propagation is a covariance matrix for each leaf yielding correlations between all landmarks involved in measurements at the same leaf.

3.6. Performance and discussion

As evident from the description in this section, there is *no approximation* involved in the update and state-recovery operations computing $\hat{x}$ from the different $p^I_{\mathbf{n}}$. The estimate $\hat{x}$ computed by this core part of treemap is the *same* as the one provided by linearized least square or EKF when using the same linearization point. Approximation errors are introduced by linearization in computing $p^I_{\mathbf{n}}$ from the original non-linear landmark–landmark constraints. When, as usual, the input are robot–landmark and robot–robot constraints, a

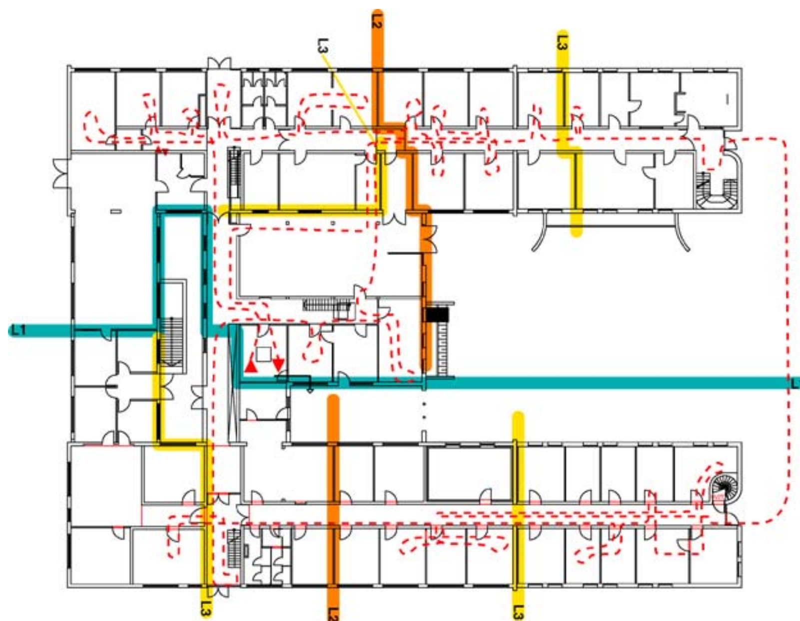

Fig. 4 DLR Institute of Robotics and Mechatronics—A typical topologically suitable building with the first three levels (L1, L2, L3) of a suitable hierarchical partitioning. It has been mapped in the experiments (Section 9), with the dashed line sketching the robots trajectory. The start and finish are indicated by small triangles

further preprocessing step is necessary (cf. Section 5). This step marginalizes out old robot poses and in doing so creates further landmark constraints. To avoid getting too many constraints, a so called sparsification is necessary, which is a second source of error. No further approximations are involved.

Local and global levels are treated conceptually the same way by least square estimation on landmarks. This is different from Atlas and the algorithm by Estrada et al. (2005), that use a graph over relations of reference frames on the global level. So as with CEKF and TJTF the division into submaps is mostly transparent for the user.

There are three key ideas that make this computation fast.

- **Many small matrices instead of one large matrix.** This is the motivation. For the matrices actually to be small the building must have a hierarchical partitioning with limited overlap (cf. Section 4) and the partitioning subalgorithm must actually find one (cf. Section 6.3).
- **Only a single path from leaf to root needs to be updated** after a new constraint is added to that leaf. Since $p^M_{\mathbf{n}}$ and $p^C_{\mathbf{n}}$ depend only on $z[\mathbf{n}{:}\downarrow]$ all other nodes still remain valid.[4] So if the tree is balanced, only $O(\log n)$ nodes are updated.
- **State-recovery is fast**, because it needs only a single matrix-vector product per node (31) to propagate the mean. Alternatively two matrix products are needed to propagate the covariance (33). This makes computing a global estimate extremely fast, because then recursive propagation is the dominant operation $O(n)$.

[4] An exception is discussed in Section 6 but does not affect the *O(log n)* claim.

Appendix A shows a worked out example for propagation of distributions in the tree corresponding to the example in Fig. 2.

4. Assumptions on topologically suitable buildings

The time needed for computation at a node $\mathbf{n}$ depends on the size of the matrices involved, which is determined by the number of landmarks in $X[\mathbf{n}{:}\downarrow\uparrow \vee \swarrow\searrow\Uparrow] = X[\mathbf{n}_\swarrow : \downarrow\uparrow \vee \mathbf{n}_\searrow : \downarrow\uparrow]$. So for each node only few landmarks should at the same time be involved in constraints below and in constraints above $\mathbf{n}$. Or intuitively speaking, the region represented by a node should only have a small border with the rest of the building.

As the experiments in Section 9 and the following considerations confirm, typical buildings allow such a hierarchical partitioning as a tree because they are hierarchical themselves, consisting of floors, corridors and rooms. Different floors are only connected through a few staircases, different corridors through a few crossings and different rooms most often only through a single door and the adjacent parts of the corridor. Thus, on the different levels of the hierarchy natural regions are: rooms, part of a corridor including adjacent rooms, one or several adjacent corridors and one or several consecutive floors (Fig. 4).

Let us formally define a *"suitable hierarchical partitioning"* and thus a *"topologically suitable building"* having such a partitioning.

Definition 1 (Suitable hierarchical partitioning). Let the measurements z be assigned to leaves of a tree. Let k be the maximum number of landmarks involved in measurements

from a single robot pose. Then the tree is a suitable hierarchical partitioning, if

1. For each node $\mathbf{n}$ the number of landmarks in $X[\mathbf{n}{:}\downarrow\uparrow]$ is $O(k)$.
2. For each leaf $\mathbf{n}$ the number of leaves $\mathbf{n}'$ for which $X[\mathbf{n}{:}\downarrow]$ and $X[\mathbf{n}' : \downarrow]$ share a landmark is $O(1)$.

Definition 2 (Topologically suitable building). A topologically suitable building is a building where a suitable hierarchical partitioning exists regardless how the robot moves.

The parameter k is small, since the robot can only observe a few landmarks simultaneously because its field of view is limited both by walls and sensor range. In particular, k does not increase when the map gets larger ($n \to \infty$). Although by this argument $k = O(1)$, the asymptotical expressions in this article explicitly show the influence of k. All expressions hold strictly if two heuristic assumptions are valid.

– The encountered building is topologically suitable, i.e. a suitable partitioning exists.
– The hierarchical tree partitioning (HTP) subalgorithm (Section 6.3) succeeds in finding such a suitable partitioning.

If Definition 1 is restricted only to leaves, it is mostly equivalent to general sparsity in the information form as exploited by SEIF and other algorithms. In general it is stronger since it demands $O(k)$ connections even between large regions. Still it is compatible with loops and nested loops as evident from the experiments (Fig. 8 contains 200 medium loops nested in 10 large loops) but it does preclude grid-like structures. So large open halls as well as most outdoor environments are not topologically suitable. If for instance an $l \times l$ square with n landmarks is divided into 2 halves, the border involves $O(l) = O(\sqrt{n})$ landmarks. So there will be an $\sqrt{n} \times \sqrt{n}$ matrix at the root node increasing update time to $O(n^{3/2})$. Estimation quality will not be affected. However, if outdoors the goal is to explore rather than to "mow the lawn" the robot will operate on a network of paths. Treemap can still be reasonable efficient then.

4.1. Computational efficiency

By Definition 1 there are $O(\frac{n}{k})$ nodes in the tree (part 2) and each stores matrices of dimension $O(k \times k)$ (part 1). Thus, the storage requirement of the treemap is $O(k^2 \cdot \frac{n}{k}) = O(nk)$ meeting requirement (R2). Updating one node takes $O(k^3)$ time for (30) and (27). State-recovery by (31) needs $O(k^2)$ time (mean only) and by (33) $O(k^3)$ time (covariance).

So after integrating new constraints into $p^I_{\mathbf{n}}$ at some leaf $\mathbf{n}$ $O(k^3 \log n)$ time is needed for updating. An estimate for the landmarks involved at some leaf $\mathbf{n}'$ can be provided in the same computation time. This way treemap can be used the same way as CEKF maintaining only local estimates but replacing CEKF's $O(kn^{3/2})$ global update with treemap's $O(k^3 \log n)$ update. As long as $\mathbf{n} = \mathbf{n}'$ we skip the treemap update and proceed as CEKF using the EKF equations in $O(k^2)$ time.

In order to compute an estimate for all landmarks, (31) must be applied recursively taking $O(k^2 \frac{n}{k}) = O(kn)$ (mean only). It will turn out in the experiments in Section 8 that the constant factor involved is extremely small. So while the possibility to perform updates in sublinear time is most appealing from a theoretical perspective, in practice treemap can compute a global estimate even for extremely large maps.

Overall, Definition 1 is both the strength and weakness of treemap. The insight that buildings have such a loosely connected topology distinguishes indoor SLAM from many other estimation problems and enables treemap's impressive efficiency. On the other hand it precludes dense planar mapping mainly ruling out outdoor environments.

5. EKF based preprocessing stage

The part of treemap discussed so far is very general. It can estimate random variables with any meaning given some Gaussian constraints with suitable topology. However it cannot marginalize out random variables that are not needed any more, i.e. old robot poses.

In this section we will derive an EKF based preprocessing stage. It receives landmark observations and odometry measurements and converts these into information on the current robot pose and information on local landmarks marginalizing out old poses. The information on landmarks is passed into the treemap as a Gaussian constraint.

In each moment there is one local region, i.e. one leaf $\mathbf{c}$ that is active corresponding to where the robot currently is. New landmark information is multiplied into $p^I_{\mathbf{c}}$ and the EKF maintains an estimate for all landmarks involved there. In this sense, the framework is similar to that of the CEKF, Atlas, and Feder's submap algorithm. Unlike Atlas and Feder's algorithm treemap employs a full least square estimator on top of this local estimate, namely the tree discussed so far. So as with CEKF, the EKF's local estimate includes information from all measurements not just from measurements in the current region.

We will first derive a simple solution where no robot pose information is passed across regions. It follows the relocation idea by Walter et al. (2005) as well as Frese and Hirzinger (2001) and sacrifices odometry information to preserve sparsity when marginalizing out old robot poses. The companion technical report (Frese, 2006b) discusses a more sophisticated sparsification scheme. While the experiments used that scheme, relocation is much easier and works very convincingly as we recently observed (Frese and Schröder, 2006).

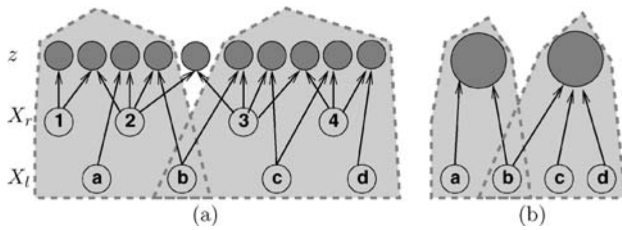

Fig. 5 *Bayesian view*. (a) A part of the example shown in (Fig. 2) with landmarks X_l (a . . . d) and robot poses X_r (1 . . . 4). The odometry constraint between pose 2 and 3 (shown outside both regions) is ignored. In this manner the robot poses are involved only in their respective region and are marginalized out. (b) The result for each region is a single Gaussian (big circle) on all landmarks in that region

5.1. Bayesian view

Figure 5 shows an example as a Bayes net with landmarks and robot poses for two regions. The odometry measurement that connects poses in both regions is ignored. Then all poses are only involved inside one region and can be marginalized out. This means, that whenever the robot enters a new region, its position is only defined by the measurements made there. The regions are however connected by overlapping landmarks. Note, though, that odometry can still be used for data association, so this does not mean the robot is actually "kidnapped". Odometry is only ignored in the sense that no constraint is integrated.

5.2. Data flow view

This process can be conveniently implemented as a preprocessing EKF (Fig. 6). When entering a region $\mathbf{c}$, treemap computes the marginal $p_{\mathbf{c}}$ (mean and covariance) of landmarks $X[\mathbf{c} : \downarrow]$ involved there.

$$p_{\text{EKF}} = p_{\mathbf{c}} = p(X[\mathbf{c}{:}\downarrow]|z_-) \tag{34}$$

We write z_- to indicate that the distribution is conditioned on the measurements made before entering $\mathbf{c}$ and z_+ for the measurements made while operating $\mathbf{c}$. The EKF is initialized with this distribution and an ∞-covariance prior for the robot pose. While the robot stays in the region, the EKF maintains

$$p'_{\text{EKF}} = p(X_r, X[\mathbf{c}{:}\downarrow]|z). \tag{35}$$

After leaving $\mathbf{c}$, information must be passed from the EKF to the treemap. For that purpose we take the EKF's marginal on landmarks

$$p(X[\mathbf{c}{:}\downarrow]|z) = \int_{X_r} p(X_r, X[\mathbf{c}{:}\downarrow]|z) \tag{36}$$

$$= \int_{X_r} p(X[\mathbf{c}{:}\downarrow]|z_-) \cdot p(X_r, X[\mathbf{c}{:}\downarrow]|z_+) \tag{37}$$

$$= p(X[\mathbf{c}{:}\downarrow]|z_-) \int_{X_r} p(X_r, X[\mathbf{c}{:}\downarrow]|z_+) \tag{38}$$

$$= p(X[\mathbf{c}{:}\downarrow]|z_-) \cdot p(X[\mathbf{c}{:}\downarrow]|z_+). \tag{39}$$

This equation relies on the odometry constraint being removed, because otherwise X_r would be involved in both factors and neither one could be moved out of the integral. The marginal is then divided ($\oslash$) by $p_{\mathbf{c}}$ the information already stored in the treemap. The result is

$$p^M_{\text{EKF}} = \frac{\int_{X_r} p'_{\text{EKF}}}{p_{\mathbf{c}}} = \frac{p(X[\mathbf{c}{:}\downarrow]|z_+) \cdot p(X[\mathbf{c}{:}\downarrow]|z_-)}{p(X[\mathbf{c}{:}\downarrow]|z_-)}$$
$$= p(X[\mathbf{c}{:}\downarrow]|z_+) \tag{40}$$

the information obtained by new measurements. It is independent from z_- and can be multiplied into $p^I_{\mathbf{c}}$ passing that information to the treemap.

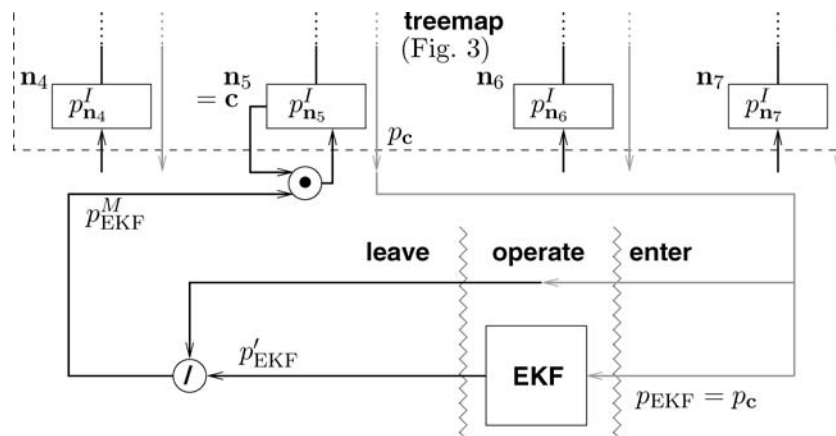

Fig. 6 *Data flow view*. The figure's top shows the lower part of the treemap (i.e. leaves) as depicted in Fig. 3. It illustrates how information is passed from the treemap into the preprocessing EKF and vice versa. When entering region $\mathbf{c}$ the EKF is initialized with the marginal $p_{\mathbf{c}}$ from the treemap. When leaving $\mathbf{c}$ again, the new information p^M_{EKF} on landmarks is multiplied into $p^I_{\mathbf{c}}$ integrating it into the treemap. Each time the robot pose is discarded and redefined by the next measurement

6. Maintenance of the tree

In this section we will discuss the bookkeeping part of the algorithm. It maintains the tree that is not defined a-priori but built while the map grows. There are three subtasks.

1. Determine for which nodes to update $p^M_{\mathbf{n}}$ and $p^C_{\mathbf{n}}$ by (24) ($\odot$), as well as (27) and (30) (Ⓜ). This task is pure bookkeeping.
2. Control the transition between the current region, $\mathbf{c}$, and the next region, $\mathbf{c}_{\text{next}}$. This defines which constraints are assigned to which leaf even though the assignment is not explicitly stored. We rely upon a heuristic that limits a region's geometric extension by *maxD*.
3. Rearrange the tree so it is balanced and well partitioned, i.e. $x[\mathbf{n}{:}\downarrow\uparrow]$ contains few landmarks in all nodes $\mathbf{n}$. Balancing is not difficult but hierarchical tree partitioning (HTP) is NP-complete. So we follow the tradition in graph partitioning (Fiduccia and Mattheyses, 1982) and optimize in greedy steps with each step being optimal.

The goal was to make treemap $O(k^3 \log n)$ in a strict asymptotical sense given that the HTP subalgorithm succeeds in finding a suitable tree (Definition 1). Unfortunately this results in a relatively involved implementation. We therefore discuss the general approach and leave the details to the pseudocode in the companion report (Frese, 2006b). In Frese and Schröder (2006) we present a simplified HTP algorithm that sacrifices the $O(k^3 \log n)$ bound.

6.1. Update

Treemap has to keep track of which landmark is involved where and when to marginalize out a landmark. So distributions $p^M_{\mathbf{n}}$, $p^C_{\mathbf{n}}$ contain a sorted list $\mathcal{L}^M_{\mathbf{n}}$, $\mathcal{L}^C_{\mathbf{n}}$ denoting the landmarks represented by the different rows/columns of the corresponding matrices and vectors. For each landmark it also contains a counter that is 1 in the leaves and added when multiplying distributions ($\odot$). There is also a global landmark array $\mathcal{L}$ with corresponding counters. We treat both as multisets writing $\uplus$ for union with adding counters and $\mathsf{l}\#\mathcal{L}$ for the counter of l in $\mathcal{L}$. Treemap detects when to marginalize out a landmark by comparing the counters passed to the node with the global counter.

$$\mathcal{L}'_{\mathbf{n}} = \mathcal{L}^M_{\mathbf{n}_\swarrow} \uplus \mathcal{L}^M_{\mathbf{n}_\searrow} \tag{41}$$

$$\mathcal{L}^M_{\mathbf{n}} = \{\mathsf{l} \in \mathcal{L}'_{\mathbf{n}} | 0 < \mathsf{l}\#\mathcal{L}'_{\mathbf{n}} < \mathsf{l}\#\mathcal{L}\} \tag{42}$$

$$\mathcal{L}^C_{\mathbf{n}} = \{\mathsf{l} \in \mathcal{L}'_{\mathbf{n}} | 0 < \mathsf{l}\#\mathcal{L}'_{\mathbf{n}} = \mathsf{l}\#\mathcal{L}\} \tag{43}$$

It further maintains an array $\mathbf{lca}[\mathsf{l}]$ storing for each landmark l the least common ancestor of all leaves involving l. It is that node, that satisfies $\mathsf{l}\#\mathcal{L}'_{\mathbf{lca}[\mathsf{l}]} = \mathsf{l}\#\mathcal{L}$ and where l is marginalized out.

If $p^I_{\mathbf{n}}$ changes—for instance by multiplying p^M_{EKF} into $p^I_{\mathbf{c}}$ – all $p^M_{\mathbf{m}}$ and $p^C_{\mathbf{m}}$ are updated from $\mathbf{m} = \mathbf{n}$ up to the root. The same applies if a new leaf has been inserted. Additionally $\mathbf{lca}[\mathsf{l}]$ can change for a landmark involved in $p^I_{\mathbf{n}}$ and all nodes from the old $\mathbf{lca}[\mathsf{l}]$ to the root are updated too. By Definition 1.2, only $O(1)$ leaves share landmarks with a given leaf so $O(\log n)$ nodes are updated in $O(k^3 \log n)$ computation time.

In the following we need to find all leaves involving a given landmark l. We recursively go down from $\mathbf{m} = \mathbf{lca}[\mathsf{l}]$ as far as $\mathsf{l} \in \mathcal{L}^M_{\mathbf{m}}$. By Definition 1.2 this holds for only $O(1)$ leaves taking $O(k \log n)$ time.

6.2. Region changing control heuristic

Let treemap currently operate in a region, i.e. a leaf $\mathbf{c}$. The EKF directly handles odometry, observation of new landmarks, and of landmarks in $X[\mathbf{c} : \downarrow]$. There are two reasons to leave $\mathbf{c}$ and enter another region $\mathbf{c}_{\text{next}}$. First a landmark may be observed that is not within $\mathbf{c}$, i.e. $X[\mathbf{c}{:}\not\downarrow]$. In this case we transition to a region containing this landmark so as to pass information on that landmark from the treemap to the EKF. A second reason is the need to limit the number of landmarks in a region for efficiency. We actually limit the distance *maxD* between landmarks in the same region instead of directly limiting the number of landmarks. This allows us to later add landmarks that have been overlooked.

As we transition from $\mathbf{c}$ to $\mathbf{c}_{\text{next}}$, the two regions must share at least two landmarks, to avoid disintegration of the map due to the omitted odometry link. Thus, treemap checks, whether $\mathbf{c}$ must be left for one of the two reasons above and determines $\mathbf{c}_{\text{next}}$ with these steps:

1. Find all leaves sharing at least two landmarks with $\mathbf{c}$.
2. For each of these leaves verify whether *maxD* would be exceeded when adding the landmarks currently in the robot's field of view.
3. Among those where it is not exceeded, choose $\mathbf{c}_{\text{next}}$ as the one that already involves most of the landmarks in the robot's field of view.
4. If all leaves would exceed *maxD* then add a new leaf as $\mathbf{c}_{\text{next}}$.
5. Leave $\mathbf{c}$. Add landmarks observed to $\mathbf{c}_{\text{next}}$ and enter $\mathbf{c}_{\text{next}}$.

When a new leaf is added, it is inserted directly above the root node. It will then be moved to a better location by the HTP subalgorithm.

6.3. Hierarchical tree partitioning (HTP)

The HTP subalgorithm optimizes the tree while the robot moves. The goal is to meet Definition 1, which is the

prerequisite for our $O(\ldots)$ analysis. The problem is equivalent to the *Hierarchical Tree Partitioning Problem* known from graph theory and parallel computing and being NP-complete. However, successful heuristic algorithms have been developed (Vijayan, 1991)—the most popular of which is the Kernighan and Lin heuristic (Fiduccia and Mattheyses, 1982). It employs a greedy strategy in each step moving that node which minimizes the cost function. Hendrickson and Leland (1995) report that it works especially well when applied hierarchically. We can do this easily since we optimize an existing tree. Overall the HTP subalgorithm makes $O(1)$ optimization steps (5 in our experiments) whenever changing regions, so the time spent in partitioning is limited. It is heuristic experience and formally part of Definition 1 that this suffices to maintain a well partitioned tree.

In each optimization step we choose one node $\mathbf{r}$ to optimize. We move a subtree somewhere left-below that node to the right side or vice versa. This affects $\mathbf{r}$ and its descendents but we only consider $\mathbf{r}$ itself, priorizing parents over children. The cost function that is optimized is

$$\mathrm{par}(\mathbf{r}) = \left|\mathcal{L}^{\mathbf{M}}_{\mathbf{r}_\swarrow}\right| + \left|\mathcal{L}^{\mathbf{M}}_{\mathbf{r}_\searrow}\right| = |\mathbf{X}[\mathbf{r}_\swarrow : \downarrow\uparrow]| + |\mathbf{X}[\mathbf{r}_\searrow : \downarrow\uparrow]| \qquad (44)$$

the number of landmarks involved in $p^M_{\mathbf{r}_\swarrow}$ and $p^M_{\mathbf{r}_\searrow}$. This number determines the size of the matrices involved in computation at $\mathbf{r}$. The subtree that we choose to move from one side of $\mathbf{r}$ to the other is that which minimizes $\mathrm{par}(\mathbf{r})$. The cost function depends only on which subtree to move, not on where to move it. Therefore the optimal subtree is found by recursively going through all descendants $\mathbf{s}$ of $\mathbf{r}$ that share a landmark with $\mathbf{r}$. At each node, $\mathrm{par}(\mathbf{r})$ is evaluated for the situation that $\mathbf{s}$ was moved to the other side of $\mathbf{r}$.

$$\begin{aligned}
\mathsf{l} \in \mathcal{L}'^M_{\mathbf{r}_\swarrow} \quad &\Leftrightarrow \quad 0 < \mathsf{l}\#\mathcal{L}^M_{\mathbf{r}_\swarrow} \mp \mathsf{l}\#\mathcal{L}^M_{\mathbf{s}} < \mathsf{l}\#\mathcal{L} \\
\mathsf{l} \in \mathcal{L}'^M_{\mathbf{r}_\searrow} \quad &\Leftrightarrow \quad 0 < \mathsf{l}\#\mathcal{L}^M_{\mathbf{r}_\searrow} \pm \mathsf{l}\#\mathcal{L}^M_{\mathbf{s}} < \mathsf{l}\#\mathcal{L} \qquad (45)
\end{aligned}$$

$$\begin{aligned}
\mathrm{par}(\mathbf{r})_{\mathbf{s}} = &\left|\left\{\mathsf{l} \middle| \mathbf{0} < \mathsf{l}\#\mathcal{L}^M_{\mathbf{r}_\swarrow} \mp \mathsf{l}\#\mathcal{L}^M_{\mathbf{s}} < \mathsf{l}\#\mathcal{L}\right\}\right| \\
&+ \left|\left\{\mathsf{l} \middle| 0 < \mathsf{l}\#\mathcal{L}^M_{\mathbf{r}_\searrow} \pm \mathsf{l}\#\mathcal{L}^M_{\mathbf{s}} < \mathsf{l}\#\mathcal{L}\right\}\right| \qquad (46)
\end{aligned}$$

The case with $-$ and $+$ corresponds to moving from $\swarrow$ to $\searrow$, $+$ and $-$ corresponds to the other way. Each evaluation is performed by (46) in $O(k)$ time using the counters in $\mathcal{L}^M_{\ldots}$. Node $\mathbf{r}$ involves $O(k)$ landmarks each in turn involved at $O(1)$ leaves so overall $O(k \log n)$ nodes are checked in $O(k^2 \log n)$ computation time. The tree should be kept balanced. Thus $\mathbf{s}$ is only considered, if after moving

$$\frac{1}{2}\mathbf{r}_{\swarrow\,\mathrm{size}} \le \mathbf{r}_{\searrow\,\mathrm{size}} \le 2\mathbf{r}_{\swarrow\,\mathrm{size}} \qquad (47)$$

where $\mathbf{n}_{\mathrm{size}}$ is the number of leaves below $\mathbf{n}$.

We still have to determine exactly where to insert $\mathbf{s}$. For $\mathrm{par}(\mathbf{n})$ it only matters, whether $\mathbf{s}$ is inserted *somewhere left-below* ($\mathrm{par}(\mathbf{n})_\swarrow$), *somewhere right-below* $\mathbf{n}$ ($\mathrm{par}(\mathbf{n})_\searrow$), or *directly above* ($\mathrm{par}(\mathbf{n}_\uparrow)_\uparrow$).

$$\mathrm{par}(\mathbf{n}_\uparrow)_\uparrow = \left|\mathcal{L}^M_{\mathbf{s}}\right| + \left|\mathcal{L}^M_{\mathbf{n}}\right| \qquad (48)$$

$$\mathrm{par}(\mathbf{n})_\swarrow = \left|\left\{\mathsf{l} \middle| 0 < \mathsf{l}\#\mathcal{L}^M_{\mathbf{n}_\swarrow} + \mathsf{l}\#\mathcal{L}^M_{\mathbf{s}} < \mathsf{l}\#\mathcal{L}\right\}\right| + \left|\mathcal{L}_{\mathbf{n}_\searrow}\right| \qquad (49)$$

$$\mathrm{par}(\mathbf{n})_\searrow = \left|\mathcal{L}_{\mathbf{n}_\swarrow}\right| + \left|\left\{\mathsf{l} \middle| 0 < \mathsf{l}\#\mathcal{L}^M_{\mathbf{n}_\searrow} + \mathsf{l}\#\mathcal{L}^M_{\mathbf{s}} < \mathsf{l}\#\mathcal{L}\right\}\right| \qquad (50)$$

So the insertion point that minimizes $\mathrm{par}(\ldots)$ priorizing parents over children can be found by descending through the tree as follows:

1. Start with $\mathbf{n} = \mathbf{r}_\searrow$ (or $\mathbf{r}_\swarrow$ resp.)
2. Evaluate $\mathrm{par}(\mathbf{n})$ for each of the three choices *directly above* (48), *somewhere left-below* (49), or *somewhere right-below* (50).
3. If directly above is best and the new parent of $\mathbf{n}$ and $\mathbf{s}$ would be balanced (47) then insert $\mathbf{s}$. Update from old and new $\mathbf{s}$ to the root.
4. Else set $\mathbf{n}$ to $\mathbf{n}_\swarrow$ or $\mathbf{n}_\searrow$ whichever is better and go to step 2.

7. Comparison with the Thin Junction Tree Filter

Paskin (2003) has proposed an algorithm, the *Thin Junction Tree Filter* (TJTF), which is closely related to the treemap algorithm, although both have been independently developed from completely different perspectives.[5] Paskin views the problem as a Gaussian graphical model. He utilizes the fact that if a set of nodes (i.e. a set of landmarks) separates the graphical model into two parts, then these parts are conditionally independent given estimates for the separating nodes. The algorithm maintains a junction tree, where each edge corresponds to such a separation, passing marginalized distributions along the edges.

Treemap's tree is very similar to TJTF's junction tree. The most important difference is how treemap and TJTF ensure that no node involves too many landmarks. TJTF further sparsifies thereby sacrificing information for computation time. Treemap on the other hand tries to rearrange the tree with its HTP subalgorithm to reduce the number of landmarks involved. It never sacrifices information except when integrating an observation into the tree the first time. There are arguments in favor of both approaches. If treemap succeeds in finding a good tree, that is certainly better than sacrificing

[5] Originally I developed treemap from a hierarchy-of-regions and linear-equation-solving perspective. I later added the Bayesian view provided in this article.

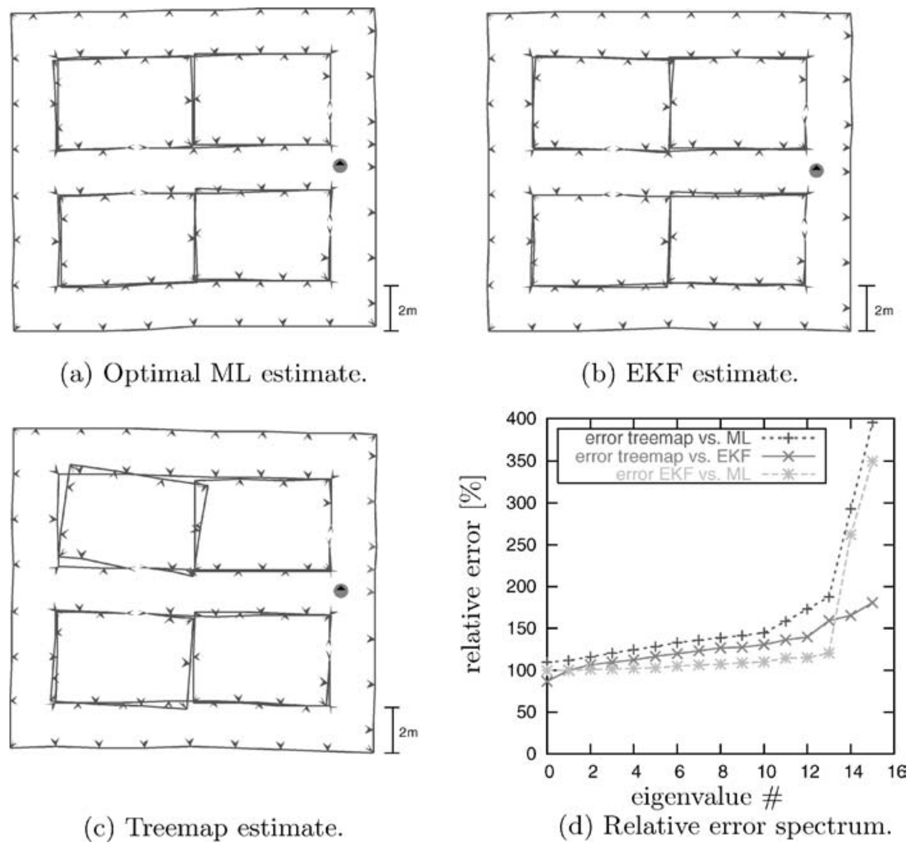

Fig. 7 Small noise simulation results. (a)–(c) shows the estimate of ML, EKF, and treemap. (d) compares the relative error as a generalized eigenvalue spectrum. Treemap performs well relative to EKF but both suffer from linearization error

information. However, if no such suitable tree exists, this question is debatable.

Consider the example in Section 4 of densely mapping an open plane. This is not topologically suitable and treemap's computation time will increase to $O(n^{3/2})$. TJTF in contrast will force each node to involve only $O(k)$ landmarks by

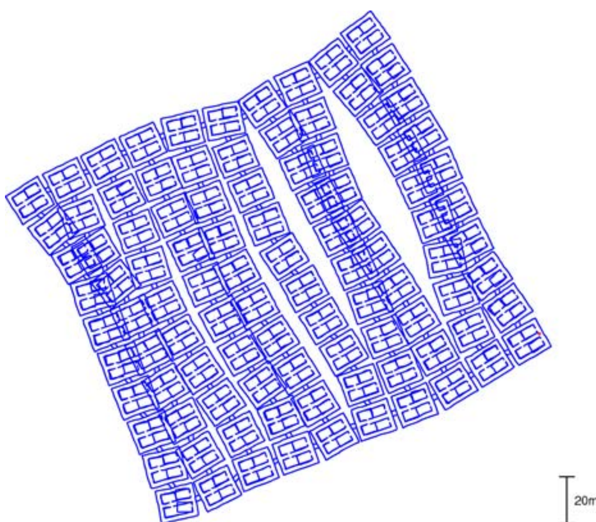

Fig. 8 Large scale simulation experiment: Treemap estimate ($n = 11300$)

sparsification, saving it's $O(k^3\ n)$ computation time. But is the posterior represented still a good approximation?

Let us consider one node of TJTF's tree that roughly divides the map into equal halves. Originally these halves have an $O(\sqrt{n})$ border where landmarks are tightly linked to both halves of the map. The considered node represents only $O(k)$ landmarks, so most of these landmarks loose their probabilistic link to one half of the map during sparsification. This is not just slightly increasing the estimation error but actually introduces breaks in the map, violating for instance (R1).

A further difference between treemap and TJTF is that treemap views its tree as a part-whole hierarchy with a designated root corresponding to the whole building. For TJTF on the other hand the tree is just an acyclic graph without designated root. This difference leads to the data flow structure in treemap whereby the posterior is updated as information matrices are passed upwards after which inference occurs with the mean and, optionally covariance calculations downwards. Together with the representation of $p^C_{\mathbf{n}}$ by H, h (27) this reduces the computation time for the mean from $O(k^3)$ per node to $O(k^2)$ per node compared to passing information matrices downwards.

Treemap saves a further factor of $O(k)$ by taking care that each landmark is only involved in $O(1)$ leaves, so there are $O(\frac{n}{k})$ nodes. This is at least typically enforced by the region

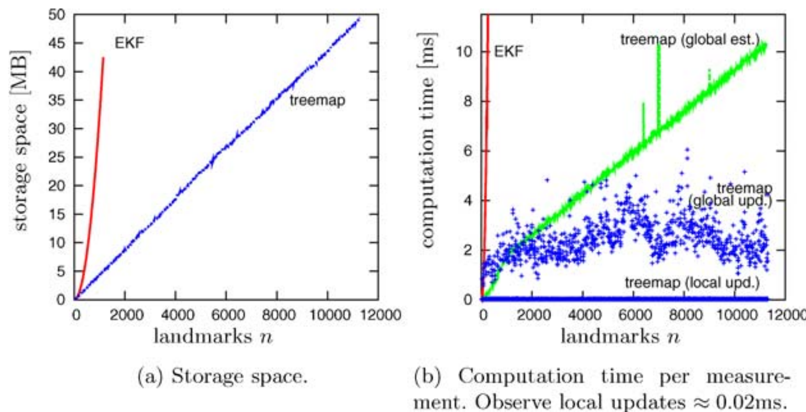

(a) Storage space.

(b) Computation time per measurement. Observe local updates ≈ 0.02ms.

Fig. 9 Large scale simulation experiment: Storage space and computation time over number of landmarks n

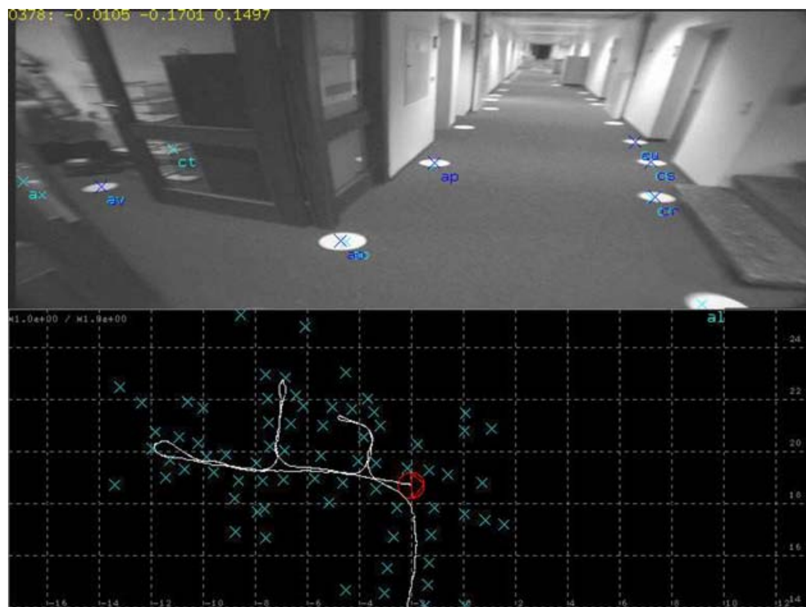

Fig. 10 Screen shot of the SLAM implementation mapping the DLR building. The corresponding video can be downloaded from the author's website: `http://www.informatik.uni-bremen.de/~ufrese/slam-videos2_e.html`

changing control heuristic, and for the analysis it is formally assumed by Definition 1.2. Thereby, treemap groups measurements as geometrically contiguous regions, whereas TJTF chooses the node that minimizes the KL divergence during sparsification. Overall this leads to a computation time for mean recovery of $O(kn)$ for treemap vs. $O(k^3 n)$ for TJTF.

Treemap maintains a balanced tree thereby limiting update and computation of a local estimate to $O(k^3 \log n)$. Paths in TJTF's tree however may have a length of $O(n)$ so it cannot update that fast exactly.

Summarizing the discussion, treemap applies a more elaborate bookkeeping to reduce computation time. This bookkeeping on the other hand makes it considerable more difficult to implement than TJTF.

8. Simulation experiments

This section presents the simulation experiments conducted to verify the algorithm with respect to the requirements (R1)–(R3). Clearly space (R2) and time (R3) consumption are straightforward to measure but how should one assess map quality with respect to requirement (R1)? It should be kept in mind, that our focus is on the core estimation algorithm, not on the overall system. So relative, not absolute, error is the quantity to be considered. This is achieved by generalized eigenvalues.

We therefore repeat the same experiment with independent measurement noise 1000 times passing the same measurements to treemap, EKF and the optimal ML estimator. We derive an error covariance matrix $C_{\text{treemap}}, C_{\text{ML}}, C_{\text{EKF}}$ for all three[6] and compare the square root of the generalized eigenvalue spectrum (Frese, 2006a). This spectrum illustrates the relative error in different aspects of the map, i.e. different linear combinations of landmark coordinates. In particular the smallest and largest eigenvalue bound the relative error of any aspect.

[6] To limit the number of necessary runs, only eight selected landmarks are used.

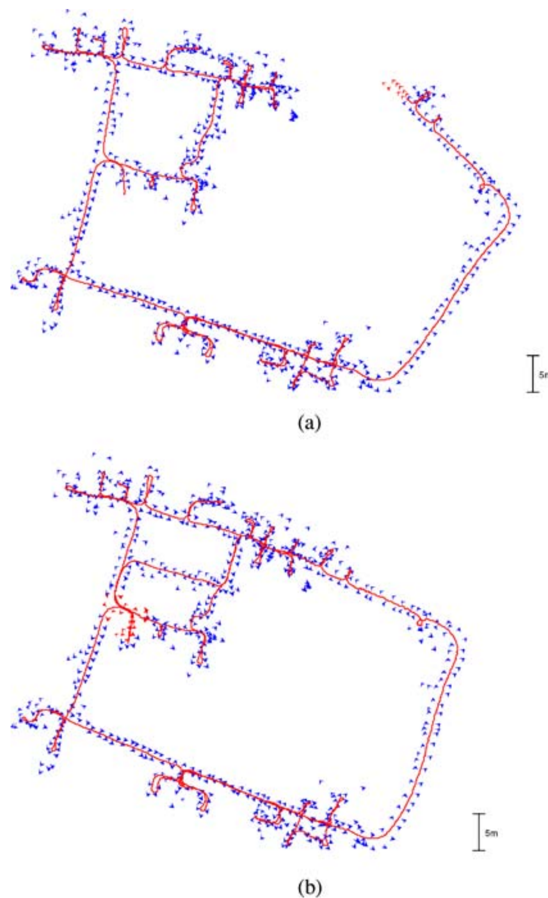

Fig. 11 (a) Treemap estimate before closing the large loop having an accumulated error of 16.18 m mainly caused by the robot leaving the building in the right upper corner. (b) Final treemap estimate after closing the large loop and returning to the starting position closing another loop

The experiments use the more sophicasted scheme for passing the robot pose described in the companion report. They have been conducted on an Intel Xeon, 2.67 GHz processor. The landmark sensor has 2.5% distance and 2° angle noise. Odometry has $0.01\sqrt{\text{m}}$ continuous velocity noise (robot radius $=$ 0.3 m). The algorithm's parameters are *optHTPSteps* = 5 optimization steps and *maxD* = 5 m region size.

8.1. Small noise experiment

The small noise simulation experiment allows statistical evaluation of the estimation error and comparison with EKF and ML (Fig. 7). At first sight all three appear to be of the same quality (except for the left upper room in the treemap estimate) and perfectly usable for navigation. The orientation of the rooms appears to be an issue. There are no overlapping landmarks between room and corridor. Thus the larger error in treemap is likely caused when changing regions.

Figure 7(d) reports the relative error as a generalized eigenvalue spectrum. When comparing treemap vs. ML, the smallest relative error is 110% (87% vs. EKF) and the largest, 395% (181% vs. EKF). The median relative error is 137% compared to ML with two outliers of 395% and 293% and median 125% compared to EKF. The outliers are also apparent in the plot comparing EKF to ML, so they are probably caused by linearization errors occurring in EKF and treemap.

8.2. Large scale map experiment

The second experiment is an extremely large map consisting of 10×10 copies of the building used before (Fig. 8). There are $n = 11300$ landmarks, $m = 312020$ measurements and $p = 63974$ robot poses. The EKF experiment was aborted due to large computation time.

In Fig. 9(a) storage space consumption is clearly shown to be linear for treemap ($O(kn)$) and quadratic ($O(n^2)$) for EKF. Overall computation time was 31.34 s for treemap and 18.89 days (extrapolated $\sim mn^2$) for EKF. Computation time per measurement is shown in Fig. 9(b). Time for three different computations is given: Local updates (dots below <0.5 ms), global updates computing a local map (scattered dots above 0.5 ms) and the additional cost for computing a global map are plotted w.r.t. n. Note that the global updates have a very fluctuating computation time because the number of nodes updated depends on the subtrees moved by the HTP subalgorithm. The spikes in the global estimation plot are caused by lost timeslices.[7]

Overall the algorithm is extremely efficient updating an $n = 11300$ landmark map in 12.37 ms. Average computation time is $1.21\,\mu\text{s} \cdot k^2$ for a local update, $0.38\,\mu\text{s} \cdot k^3 \cdot \log n$ for a global update, and $0.15\,\mu\text{s} \cdot kn$ for a global map (mean only), respectively ($k \approx 5.81$). The latter is surely the most impressive practical result. It allows the computation of a global map even for extremely large n, avoiding the complications of local map handling. Recently, we could even improve this result updating an $n = 1033009$ landmarks map in 443 ms or in 23 ms for a local update of ≈ 10000 landmarks (Frese and Schröder, 2006).

9. Real world experiments

The real world experiment reported in this section shows how treemap works in practice by mapping the DLR Institute of Robotics and Mechatronics' building (Fig. 4) that serves as an example of a typical office building. The robot is equipped with a camera system (field of view: $\pm 45°$) at a height of 1.55 m and controlled manually. We

[7] Processing time had 5 ms resolution, so clock time has been used.

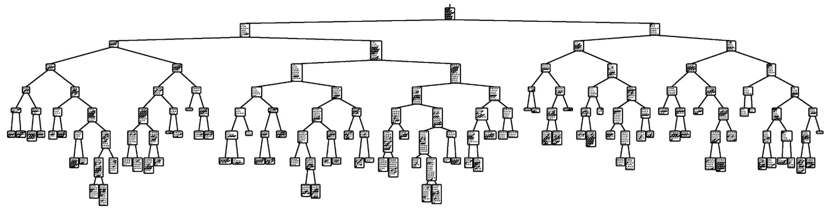

Fig. 12 Tree representation of the map. The size of the node ovals is proportional to number of landmarks represented

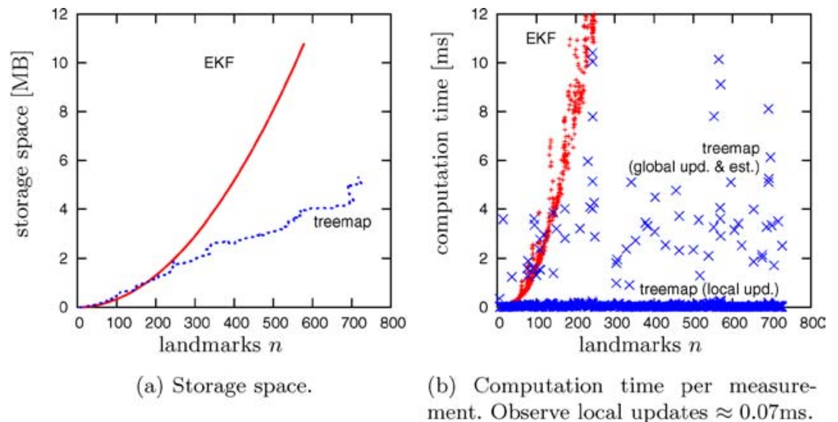

(a) Storage space.

(b) Computation time per measurement. Observe local updates $\approx$ 0.07ms.

Fig. 13 Real experiment performance

set circular fiducials throughout the floor (Fig. 10) that were visually detected by Hough-transform and a gray-level variance criterion (Otsu, 1979).

Since the landmarks are identical, identification is based on their relative position employing two different strategies in parallel. Local identification is performed by simultaneously matching all observations from a single robot pose to the map, taking into account both error in each landmark observation and error in the robot pose. For global identification we encountered considerable difficulties in detecting closure of a loop. Before closing the largest loop, the accumulated robot pose error was 16.18 m (Fig. 11) and the average distance between adjacent landmarks was ≈ 1 m. With indistinguishable landmarks, matching observations from a single image was not reliable enough.

Instead, the algorithm matches a map patch of radius 5 m around the robot. When the map patch is recognized somewhere else in the map, the identity of all landmarks in the patch is changed accordingly and the loop is closed (Frese, 2004). It is a particular advantage of the treemap algorithm to be able to change the identity of landmarks already integrated into the map. This allows the use of the *lazy data association* framework by Hähnel et al. (2003). The regions were $maxD = 7$ m large.

The final map contains $n = 725$ landmarks, $m = 29142$ measurements and $p = 3297$ robot poses (Fig. 11). The results highlight the advantage of using SLAM, because after closing the loop the map is much better. Figure 12 shows the internal tree representation ($k \approx 16.39$). The tree is balanced and well partitioned, i.e. no node represents too many landmarks. It can be concluded that the building is indeed topologically suitable in the sense discussed in Section 4.

If only a local update is performed, as is often the case (Fig. 13), then computation time is extremely low. The average time is 0.77 μs$\cdot k^2$ for a local update, 0.02μs$\cdot k^3 \log n$ for a global update and 0.04 μs $\cdot\, kn$ for a global map (mean) respectively. Accumulated computation time is 2.95 s for treemap and 660 s (extrapolated $\sim mn^2$) for EKF.

10. Conclusion

The treemap SLAM algorithm proposed in this article works by dividing the map into a hierarchy of regions represented as a binary tree. With this data structure, the computations necessary for integrating a measurement are limited essentially to updating a leaf of the tree and all its ancestors up to the root. From a theoretical perspective the main advantage is that a local map can be computed in $O(k^3 \log n)$ time. In practice, it is equally important that a global map can be computed in $O(kn)$ time allowing the update of a map with $n = 11300$ landmarks in 12.37 ms on an Intel Xeon, 2.67 GHz. Treemap is exact up to linearization and sparsification where some information on the landmarks is sacrificed to pass information on the robot pose between regions. Still despite the sparsification, if two landmarks are observed together, the fact that we know their precise relative location will be reflected by the estimate after the next update.

With respect to the three criteria (R1)–(R3) proposed in Section 1, the algorithm was verified theoretically, by simulation experiments, and by experiments with a real robot.

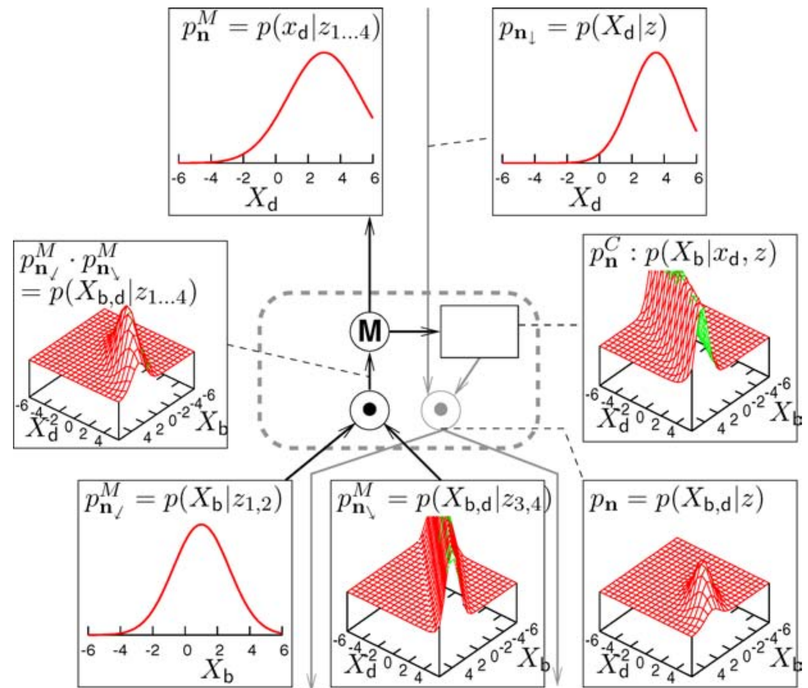

Fig. 14 The probability distributions involved in computation at node **n** in Fig. 2. The example assumes 1D landmarks. Constraint z_1 declares X_a to be 0 with variance 2. Constraints $z_{2...7}$ declare the difference between successive landmarks to be 1 with variance 1. With only these constraints the result would be $E(X \mid z_{1...7}) = (0\ 1\ 2\ 3\ 4\ 5\ 6)^T$. The last constraint z_8 contradicts by declaring X_g as 7 with variance 2. Thus the estimate stretches to $E(X \mid z_{1...8}) = (0.2\ 1.3\ 2.4\ 3.5\ 4.6\ 5.7\ 6.8)^T$. The text explains how treemap computes this result

There are two preconditions for achieving these results. (a) The environment must be topologically suitable, i.e. have a hierarchical partitioning and (b) the HTP subalgorithm must find one as explained in Section 4. This is indeed a major drawback since it precludes mapping dense outdoor environments. The second drawback is that performing the bookkeeping in $O(k^3 \log n)$ significantly complicates the algorithm. Consequently we have simplified the algorithm meanwhile (Frese and Schröder, 2006). We plan to generalize it so it can handle different SLAM variants such as 2D-, 3D-, landmark-, pose-, and bearing-only-SLAM and publish the code as an open source implementation.

A major future challenge will be uncertain data-association. Some authors address this issue with a framework that evaluates the likelihood of several data association hypotheses (Hähnel et al., 2003). Regardless of how the overall framework is conceived, it needs a classical SLAM algorithm as the core engine to evaluate a single hypothesis. Efficiency is then even more crucial because updates must be performed for each of the hypothesis considered.

Appendix

A. Example for propagation of distributions in the tree

Figure 14 shows the computation at node **n** in the example from Fig. 2. Gaussians are denoted by $\mathcal{N}(\mu, C)$ in covariance form and by $\mathcal{N}^{-1}(b, A)$ in information form. On the left $z[\mathbf{n}_\swarrow : \downarrow] = z_{1,2}$, $X[\mathbf{n}_\swarrow : \downarrow] = X_{a,b}$ but $X[\mathbf{n}_\swarrow : \downarrow\uparrow] = X_b$. So $\mathbf{n}_\swarrow$ passes

$$p^M_{\mathbf{n}_\swarrow} = p(X_b|z_{1,2}) = \mathcal{N}^{-1}\left(-\frac{2}{3}, \frac{1}{3}\right) = \mathcal{N}(1, 3) \tag{51}$$

to **n**. On the right $z[\mathbf{n}_\searrow : \downarrow] = z_{3,4}$, $z[\mathbf{n}_\searrow : \downarrow] = X_{b,c,d}$, and $X[\mathbf{n}_\searrow : \downarrow\uparrow]$ consists of X_b (shared with $\mathbf{n}_\swarrow$) and X_d (shared with $\mathbf{n}_\uparrow$). So $\mathbf{n}_\searrow$ passes

$$p^M_{\mathbf{n}_\searrow} = p(X_{b,d}|z_{3,4}) = \mathcal{N}^{-1}\left(\begin{pmatrix} 2 \\ -2 \end{pmatrix}, \begin{pmatrix} 1/2 & -1/2 \\ -1/2 & 1/2 \end{pmatrix}\right). \tag{52}$$

to **n**. This Gaussian is degenerate, because $z_{3,4}$ contain no absolute position information. It declares the difference $X_d - X_b$ as $\mathcal{N}(2, 2)$ plus an infinite uncertainty for $X_d + X_b$. This is typical for SLAM and no problem in the information form. Node **n** multiplies both ($\odot$) by (24).

$$\begin{aligned} p^M_{\mathbf{n}_\swarrow} \cdot p^M_{\mathbf{n}_\searrow} &= p(X_{b,d}|z_{1...4}) \\ &= \mathcal{N}^{-1}\left(\begin{pmatrix} 4/3 \\ -2 \end{pmatrix}, \begin{pmatrix} 5/6 & -1/2 \\ -1/2 & 1/2 \end{pmatrix}\right) \\ &= \mathcal{N}\left(\begin{pmatrix} 1 \\ 3 \end{pmatrix}, \begin{pmatrix} 3 & 3 \\ 3 & 5 \end{pmatrix}\right) \end{aligned}$$

The next step is to marginalize (Ⓜ) out $X[\mathbf{n}{:}\swarrow\searrow\updownarrow] = X_b$ by (28).

$$p^M_{\mathbf{n}} = p(X_d|z_{1...4}) = \mathcal{N}^{-1}(-6/5, 1/5) = \mathcal{N}(3, 5) \tag{53}$$

$$p^C_{\mathbf{n}} = p(X_b|x_d, z) = \mathcal{N}(3/5\, x_d - 4/5, 1) \tag{54}$$

Note, that the mean of $p^C_{\mathbf{n}}$ is defined as a linear function $Hx_{\mathsf{d}} + h$ of x_{d}. In this example $H = (3/5)$ and $h = (-4/5)$ by (27). H, h and the covariance $P^{-1} = (6/5)$ are stored at $\mathbf{n}$. $p^M_{\mathbf{n}}$ is passed to $\mathbf{n}_\uparrow$. Up to now only $z_{1...4}$ have been integrated. Imagine $p^M_{\mathbf{n}}$ was directly fed back into $\mathbf{n}$ using it as $p_{\mathbf{n}\uparrow}$, which should actually be computed by the parent. Then $z_{5...8}$ were bypassed and $\mathbf{n}$ would provide $\hat{x}_{\mathsf{a...d}} = (\,0\;1\;2\;3\,)$. Instead at some point above $\mathbf{n}$ the contradicting distributions $p(X_{\mathsf{d}}|z_{1...4}) = \mathcal{N}(3, 5)$ and $p(X_{\mathsf{d}}|z_{5...8}) = \mathcal{N}(4, 5)$ are integrated and

$$p_{\mathbf{n}_\uparrow} = p(X_{\mathsf{d}}|z) = \mathcal{N}(3.5, 5/2) \tag{55}$$

is passed from $\mathbf{n}_\uparrow$ to $\mathbf{n}$. It is multiplied ($\odot$) with $p^C_{\mathbf{n}}$ by (33) yielding

$$p_{\mathbf{n}} = p(X_{\mathsf{b,d}}|z) = \mathcal{N}\left(\begin{pmatrix} 1.3 \\ 3.5 \end{pmatrix}, \begin{pmatrix} 21/10 & 3/2 \\ 3/2 & 5/2 \end{pmatrix}\right) \tag{56}$$

which is passed down to $\mathbf{n}_\swarrow$ and $\mathbf{n}_\searrow$.

Acknowledgments I would like to thank Bernd Krieg-Brückner, Robert Ross, Matthew Walter, and the anonymous reviewers for valuable comments.

References

Bosse, M., Newman, P., Leonard, J., and Teller, S. 2004. SLAM in large-scale cyclic environments using the atlas framework. *International Journal on Robotics Research*, 23(12):1113–1140.

Duckett, T., Marsland, S., and Shapiro, J. 2000. Learning globally consistent maps by relaxation. In *Proceedings of the IEEE International Conference on Robotics and Automation*. San Francisco, pp. 3841–3846.

Duckett, T., Marsland, S., and Shapiro, J. 2002. Fast, on-line learning of globally consistent maps. *Autonomous Robots*, 12(3):287–300.

Eliazar, A. and Parr, R. 2003. DP-SLAM: Fast, robust simulataneous localization and mapping without predetermined landmarks. In *Proceedings of the International Joint Conference on Artificial Intelligence*. Acapulco, pp. 1135–1142.

Estrada, C., Neira, J., and Tardós, J. 2005. Hierarchical SLAM: Real-time accurate apping of large environments. *IEEE Transactions on Robotics*, 21(4):588–596.

Fiduccia, C. and Mattheyses R. 1982. A linear-time heuristic for improving network partitions. In *Proceedings of the 19th ACM/IEEE Design Automation Conference, Las Vegas*. pp. 175–181.

Frese, U. 2004. An $O(\log n)$ Algorithm for simulateneous localization and mapping of mobile robots in indoor environments. Ph.D. thesis, University of Erlangen- Nürnberg. `http://www.opus.ub.uni-erlangen.de/opus/volltexte/2004/70/`.

Frese, U. 2005. A proof for the approximate sparsity of SLAM information matrices. In *Proceedings of the IEEE International Conference on Robotics and Automation,* Barcelona, pp. 331–337.

Frese, U. 2006a. A discussion of simultaneous localization and mapping. *Autonomous Robots*, 20(1):25–42.

Frese, U. 2006b. Treemap: An $O(\log n)$ algorithm for indoor simultaneous localization and mapping. Technical Report 006-03/2006, Universität Bremen, SFB/TR 8 Spatial Cognition.

Frese, U. and Hirzinger, G. 2001. Simultaneous localization and mapping—A discussion. In *Proceedings of the IJCAI Workshop on Reasoning with Uncertainty in Robotics,* Seattle, pp. 17–26.

Frese, U., Larsson, P., and Duckett, T. 2004. A multigrid algorithm for simultaneous localization and mapping. *IEEE Transactions on Robotics*, 21(2):1–12.

Frese, U. and Schröder, L. 2006. Closing a million-landmarks loop. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems,* Beijing. submitted.

Gauss, C. 1821. Theoria combinationis observationum erroribus minimis obnoxiae. *Commentationes societatis regiae scientiarum Gottingensis recentiores*, 5:6–93.

Guivant, J. and Nebot, E. 2001. Optimization of the simultaneous localization and map-building algorithm for real-time implementation. *IEEE Transactions on Robotics and Automation*, 17(3):242–257.

Guivant, J. and Nebot, E. 2003. Solving computational and memory requirements of feature-based simultaneous localization and mapping algorithms. *IEEE Transactions on Robotics and Automation* 19(4):749–755.

Hähnel, D., Burgard, W., Wegbreit, B., and Thrun, S. 2003. Towards lazy data association in SLAM. In *Proceedings of the 10th International Symposium of Robotics Research.*

Hendrickson, B. and Leland, R. 1995. A multilevel algorithm for partitioning graphs. In *Proceedings of the ACM International Conference on Supercomputing,* Sorrento, pp. 626–657.

Horn, R. and Johnson, C. 1990. *Matrix Analysis*. Cambridge University Press.

Leonard, J. and Feder, H. 2001. Decoupled stochastic mapping. *IEEE Journal of Ocean Engineering*, 26(4):561–571.

Montemerlo, M., Thrun, S., Koller, D., and Wegbreit, B. 2002. FastSLAM: A factored solution to the simultaneous localization and mapping problem. In *Proceedings of the AAAI National Conference on Artificial Intelligence,* Edmonton, pp. 593–598.

Nieto, J., Guivant, J., Nebot, E., and Thrun, S. 2003. Real time data association for fastSLAM. In *Proceedings of the IEEE Conference on Robotics and Autonomation,* Taipeh, pp. 412–418.

Otsu, N. 1979. A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man and Cybernetics*, 9(1):62–66.

Paskin, M. 2003. Thin junction tree filters for simultaneous localization and mapping. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, San Francisco, pp. 1157–1164.

Press, W., Teukolsky, S., Vetterling, W., and Flannery, B. 1992. *Numerical Recipes*, 2nd edn. Cambridge University Press, Cambridge.

Smith, R., Self, M., and Cheeseman, P. 1988. Estimating uncertain spatial relationships in robotics. In I Cox, and G. Wilfong (eds), *Autonomous Robot Vehicles*. Springer Verlag, New York, pp. 167–193.

Stachniss, C. and Burgard, W. 2004. Exploration with active loop-closing for fastSLAM. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1505–1510.

Thrun, S., Burgard, W., and Fox, D. 2005. *Probabilistic Robotics*. MIT Press.

Thrun, S., Liu, Y., Koller, D., Ng, A., Ghahramani, Z., and Durrant-Whyte, H. 2004. Simultaneous localization and mapping with sparse extended information filters. *International Journal of Robotics Research*, 23(7–8):613–716.

Vijayan, G. 1991. Generalization of min-cut partitioning to tree structures and its applications. *IEEE Transactions on Computers*, 40(3):307–314.

Walter, M., Eustice, R., and Leonard, J. 2005. A provably consistent method for imposing exact sparsity in feature-based SLAM information filters. In *Proceedings of the 12th International Symposium of Robotics Research.*

Udo Frese was born in Minden, Germany in 1972. He received the Diploma degree in computer science from the University of Paderborn in 1997. From 1998 to 2003 he was a Ph.D. student at the German Aerospace Center in Oberpfaffenhofen. In 2004 he received his Ph.D. degree from University of Erlangen-Nürnberg and joined SFB/TR 8 Spatial Cognition at University of Bremen. He works on mobile robotics, SLAM and computer vision.