

# A Multigrid Approach for Accelerating Relaxation-based SLAM

Udo Frese

Tom Duckett

**DLR Oberpfaffenhofen  
German Aerospace Center  
Institute of Robotics und Mechatronics  
D-82230 Wessling, Germany**

**AASS Learning Systems Lab  
Department of Technology  
Örebro University  
S-70182 Örebro, Sweden**

Udo.Frese@dlr.de, www.robotic.dlr.de    tdt@tech.oru.se, www.aass.oru.se

## Abstract

This paper addresses the problem of simultaneous localisation and mapping (SLAM) by a mobile robot. An incremental SLAM algorithm is introduced that is derived from so-called multigrid methods used for solving partial differential equations. The approach overcomes the relatively slow convergence of previous relaxation methods because it optimizes the map at multiple levels of resolution. The resulting algorithm has an update time that is linear in the number of mapped features, even when closing very large loops, and offers advantages in handling non-linearities compared to previous approaches. Experimental comparisons with alternative algorithms using two well-known data sets are also presented.

## 1 Introduction

To navigate in unknown environments, an autonomous robot requires the ability to build its own map while maintaining an estimate of its own position. The SLAM problem is hard because the same sensor data must be used for both mapping and localisation. We can separate two major sources of uncertainty in solving this problem: (*i.*) the *continuous* uncertainty in the positions of the robot and observed environmental features, and (*ii.*) the *discrete* uncertainty in the identification and re-identification of environmental features (data association). Any approach to the SLAM problem that considers both types of uncertainty must somehow search the space of possible *maps*, since alternative assignments in data association can produce very different maps.

Our approach belongs to a family of techniques where the environments is represented by a graph of spatial relations between reference frames that is obtained by scan matching [Lu and Milios, 1997; Gutmann and Konolige, 1999]. With this approach, it is natural to separate the topological (discrete) and geometric (continuous) elements of the representation, and to consider tracking the  $M$  most likely topological hypotheses as a practical so-

lution to the SLAM problem. Alternative topological hypotheses generally correspond to decisions over whether or not to “close a loop”, based on the uncertainty in the re-identification of previously mapped features. The key problem here is that to evaluate the likelihood of one single hypothesis, a large linear equation system has to be solved in order to infer the most likely geometric representation given a particular topology.

A desirable property for any SLAM algorithm is that the computation time for updates should be linear in the number of features  $n$  stored in the map [Frese and Hirzinger, 2001]. To achieve this objective, we have investigated so-called multigrid methods for solving partial differential equations [Briggs, 1999], resulting in a new SLAM algorithm for solving the equation system of a single topological hypothesis called ‘Multilevel Relaxation’. Relaxation is an iterative method for solving equations, which is equivalent to Gauss-Seidel iteration or Gibbs sampling at zero temperature. The new approach improves on the previously introduced algorithm [Duckett *et al.*, 2002], which was significantly slower when closing large loops, by carrying out the optimization process at multiple levels of resolution in the underlying map.

In the following sections, we derive the basic algorithm for single level relaxation (§2), followed by an overview of multigrid methods (§3) and details of the Multilevel Relaxation algorithm (§4). Results including experimental comparisons with alternative algorithms are presented in §5, and the conclusion in §6 discusses how to embed the new algorithm within a framework for tracking multiple topological hypotheses.

### 1.1 Related Work

Guivant and Nebot [2001] introduced an extended Kalman filter (EKF) called Compressed SLAM (CSLAM) for real-time mapping. By restricting the Kalman update to a subset of landmarks in a local area, updates can be performed at cost  $O(1)$  and then transferred to the overall map in  $O(n^2)$ . With a further approximation, that can be reduced to  $O(n)$ , though the problem of closing large loops is not yet solved.

Montemerlo et al. [2002] used a particle filter to track the pose of the robot, where each particle also includes a set of Kalman filters estimating the position for each landmark. This approach is able to represent and search between multiple hypotheses for the full map (i.e., robot pose plus all landmark positions), but the particle set must be large enough to contain a particle sufficiently close to the true pose of the robot at all times. The algorithm requires  $O(M \log n)$  time for  $M$  particles, though it is not clear how the number of particles scales with the complexity of the environment.

Thrun et al. [2002] applied extended information filters utilizing the sparsity of the information matrices in SLAM, as proposed by Frese and Hirzinger [2001]. The equation solving is performed iteratively by relaxation. The authors propose to relax only  $O(1)$  landmarks at each step, which would result in a constant time algorithm. However, in the numerical literature, relaxation is reputed to need  $O(n^2)$  time for reducing the equation error by a constant factor [Briggs, 1999; Press *et al.*, 1992, §19.5]. For instance after observing  $n$  landmarks each  $O(1)$  times, the algorithm will have spent only  $O(n)$  time on equation solving, so it is doubtful whether this approach will suffice in general.

## 2 Single Level Relaxation

The input to the algorithm is a set  $\mathcal{R}$  of  $m = |\mathcal{R}|$  relations on  $n$  planar frames. Each relation  $r \in \mathcal{R}$  describes the likelihood distribution of frame  $a^r$  relative to frame  $b^r$ . It is modelled as a Gaussian with mean  $\mu^r$  and covariance  $C^r$ . The output is the maximum likelihood (ML) estimation vector  $\hat{x}$  for the poses of all the frames.

In the context of SLAM, each frame corresponds to the robot pose at a certain time. Each relation corresponds to a measurement of the relative pose between two frames, either by odometry for consecutive frames or as the result of matching the laser scans (or other sensor readings) taken at the respective robot poses. As usual, the mean  $\mu^r$  of such a relation is the actual measurement and the covariance  $C^r$  is taken from a suitable model of the measurement uncertainty.

The algorithm proceeds in three steps [Press *et al.*, 1992, §15]:

1. Linearize the measurement functions.
2. Compute a quadratic error function  $\chi^2(x)$  and represent it by a matrix  $A$  and a vector  $b$  as  $\chi^2(x) = x^T A x - 2x^T b$ .
3. Find the minimum  $\hat{x}$  of  $\chi^2(x)$  by solving  $Ax = b$ .

The first two steps are the same used in most least square nonlinear model fitting algorithms. Specific to relaxation is the way of solving  $Ax = b$ . It is performed by going through all block rows  $A_i$  and solving  $(Ax)_i = b_i$  for  $x_i$ . This process is repeated until convergence.

## 2.1 Derivation of the Linear Equation System

Maximizing likelihood is equivalent to minimizing negative log likelihood or  $\chi^2$  error energy:

$$\chi^2(x) = \sum_{r \in \mathcal{R}} z^r T (C^r)^{-1} z^r, \quad (1)$$

$$\text{with } z^r = f(x_{a^r}, x_{b^r}) - \mu^r, \quad (2)$$

$$f\left(\begin{pmatrix} a_x \\ a_y \\ a_\phi \end{pmatrix}, \begin{pmatrix} b_x \\ b_y \\ b_\phi \end{pmatrix}\right) = \begin{pmatrix} (a_x - b_x) \cos b_\phi + (a_y - b_y) \sin b_\phi \\ -(a_x - b_x) \sin b_\phi + (a_y - b_y) \cos b_\phi \\ a_\phi - b_\phi \end{pmatrix}$$

The measurement function  $f$  maps the two poses of the two frames  $a^r$  and  $b^r$  to the relative pose of  $a^r$  with respect to  $b^r$ . As usual, it is linearized at some linearization point  $\check{a}^r, \check{b}^r$  corresponding to some estimate for the two frames. We use the most recent estimate for  $\check{b}^r$  and choose  $\check{a}^r$  so that  $f(\check{a}^r, \check{b}^r) = \mu^r$ . This means that the linearization points chosen for a measurement are consistent with the measurement itself. Compared to using the most recent estimate for  $\check{a}^r$  this produces a much smaller error when closing a loop.

$$f\left(\begin{pmatrix} a_x \\ a_y \\ a_\phi \end{pmatrix}, \begin{pmatrix} b_x \\ b_y \\ b_\phi \end{pmatrix}\right) \approx \mu + \underbrace{\begin{pmatrix} \cos \check{b}_\phi & \sin \check{b}_\phi & 0 \\ -\sin \check{b}_\phi & \cos \check{b}_\phi & 0 \\ 0 & 0 & 1 \end{pmatrix}}_{J_a} (a - \check{a}) + \underbrace{\begin{pmatrix} -\cos \check{b}_\phi & -\sin \check{b}_\phi & \mu_y^r \\ \sin \check{b}_\phi & -\cos \check{b}_\phi & -\mu_x^r \\ 0 & 0 & -1 \end{pmatrix}}_{J_b} (b - \check{b})$$

The linearized measurement function is substituted into (2) yielding a quadratic approximation:

$$z^r \approx J_a^r (x_{a^r} - \check{a}^r) + J_b^r (x_{b^r} - \check{b}^r) \quad (3)$$

$$\chi^2(x) \approx \text{const}$$

$$+ \sum_{r \in \mathcal{R}} \begin{pmatrix} x_{a^r}^T J_a^{rT} (C^r)^{-1} J_a^r x_{a^r} \\ + x_{a^r}^T J_a^{rT} (C^r)^{-1} J_b^r x_{b^r} \\ + x_{b^r}^T J_b^{rT} (C^r)^{-1} J_a^r x_{a^r} \\ + x_{b^r}^T J_b^{rT} (C^r)^{-1} J_b^r x_{b^r} \\ - 2 x_{a^r}^T J_a^{rT} (C^r)^{-1} (J_a^r \check{a}^r + J_b^r \check{b}^r) \\ - 2 x_{b^r}^T J_b^{rT} (C^r)^{-1} (J_a^r \check{a}^r + J_b^r \check{b}^r) \end{pmatrix} \quad (4)$$

The terms involve  $x$  either quadratically (lines 1–4) or linearly (lines 5–6). They can be sorted by rows of  $x$  (either  $a^r$  or  $b^r$ ) and grouped into matrix  $A$  and vector  $b$ :

$$= x^T \underbrace{\sum_{r \in \mathcal{R}} \begin{pmatrix} \dots & J_a^{rT} (C^r)^{-1} J_a^r & \dots & J_a^{rT} (C^r)^{-1} J_b^r & \dots \\ \dots & J_b^{rT} (C^r)^{-1} J_a^r & \dots & J_b^{rT} (C^r)^{-1} J_b^r & \dots \end{pmatrix}}_A x - 2x^T \underbrace{\sum_{r \in \mathcal{R}} \begin{pmatrix} J_a^{rT} (C^r)^{-1} (J_a^r \check{a}^r + J_b^r \check{b}^r) \\ J_b^{rT} (C^r)^{-1} (J_a^r \check{a}^r + J_b^r \check{b}^r) \end{pmatrix}}_b \quad (5)$$

Each relation  $r$  contributes to block-rows  $a^r$  and  $b^r$  of  $b$  and the intersection of these rows and columns in  $A$ .

FOR Level $h$ from fine to coarse	
Relax equation on level $h$ : $A^h x^h = b^h$	
Restrict residual to next level $H$ : $b^H = I_h^H (b^h - A^h x^h)$	
Solve equation on coarsest level $A^H x^H = b^H$	
FOR Level $H$ from coarse to fine	
Interpolate solution to next level $h$ : $x^h = x^H + I_H^h x^H$	
Relax equation on level $h$ : $A^h x^h = b^h$	

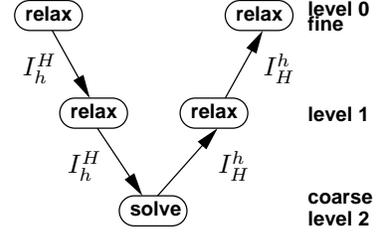


Figure 1: General multigrid algorithm (V-cycle), and example with 3 levels.

Since  $\chi^2$  is invariant under movement of the whole map,  $A$  is singular. To make it positive definite, a relation between frame 0 and a global frame is added ( $J_b = 0$ ).

The matrix  $A$  is called the information matrix, and is the inverse of the estimation covariance matrix. A block  $A_{ab} \neq 0$  appears only between frames  $a, b$  with a common relation, which are normally only  $O(1)$  for a given  $a$ . This sparsity is essential for the efficiency of relaxation. The ML estimate  $\hat{x}$  minimizes  $\chi^2(x)$  or equivalently makes the gradient equal to 0:

$$0 = \frac{\partial (\chi^2(x))}{\partial x} = \frac{\partial (x^T A x - 2x^T b)}{\partial x} = A x - b \quad (6)$$

So with the definitions made above, the equation to be solved is  $Ax = b$  for a sparse matrix  $A$ . The full posterior distribution is in principle given by  $\exp(-\frac{1}{2}\chi^2(x))$ , but this is of little practical value, since usually the posterior for some selected frames is desired. This requires computation of the covariance matrix  $A^{-1}$ , which appears to be impossible in less than  $O(n^2)$  time, so in this paper we concentrate on computing the ML estimate  $\hat{x}$ .

## 2.2 Iterative Solution by Relaxation

The basic idea of relaxation is to solve the equation system  $Ax = b$  one (block-) row at a time. Relaxation of (block-) variable  $x_i$  consists of solving (block-) row  $i$  of the equation for  $x_i$  considering all other  $x_j$  as fixed<sup>1</sup>:

$$x'_i = x_i + A_{ii}^{-1} (b_i - A_{i\bullet} x) \quad (7)$$

From the perspective of minimizing  $x^T A x - 2x^T b$ , this means finding the minimum  $x_i$  if all other  $x_j$  remain unchanged. In a single iteration, (7) is used to update all  $x_i$ . After  $x_i$  is updated, the new value is used in the update of all following  $x_j, j > i$  (Gauss-Seidel relaxation).

Every iteration reduces  $x^T A x - 2x^T b$ , so it will converge to the unique minimum  $A^{-1}b$ , thereby solving the equation. Since  $A$  is sparse, evaluating (7) takes  $O(1)$  and a single iteration  $O(n)$  time. For typical  $A$ ,  $O(n)$  iterations are needed to reduce the error by a constant factor [Briggs, 1999; Press *et al.*, 1992, §19.5]. However local or oscillating parts of the error are reduced much more effectively than smooth or global parts, so in practice, few (1 – 3) iterations suffice, except when closing a large loop [Duckett *et al.*, 2002].

<sup>1</sup> $A_{i\bullet}$  denotes row  $i$  and  $A_{\bullet i}$  denotes column  $i$  of  $A$ .

## 3 Multigrid Linear Equation Solvers

Historically, relaxation has been widely used for the numerical solution of partial differential equations (PDE). These continuous equations appear, for instance, in the simulation of heat flow, fluid dynamics or structural mechanics. As an example, the solution to a heat flow problem is a function  $\mathbb{R}^3 \rightarrow \mathbb{R}$  assigning a temperature to each point in 3-D space. Numerically they are solved by discretizing the function onto a grid of sampling points. Thereby the PDE is converted into an ordinary sparse linear equation system. It is often solved using relaxation. The problem with this approach is that *oscillating* parts of the error are reduced efficiently, but it takes much longer to reduce the remaining *smooth* error.

A breakthrough was the development of so-called multigrid methods in the 1970's [Brandt, 1977; Briggs, 1999]. The idea is to discretize the PDE at different levels of resolution. Relaxation on a fine level (high resolution) effectively smooths the error. Then relaxation on a coarser level is used to reduce that error, which on the lower resolution is again more oscillatory.

### 3.1 Geometric Multigrid

To realize this idea, a single iteration of relaxation is first performed at the finest level. The remaining residual  $b^h - A^h x^h$  is then restricted to the next coarser level by a *restriction* operator  $I_h^H$ <sup>2</sup>. On the coarsest level, the residual equation is solved directly (e.g., by Cholesky decomposition [Press *et al.*, 1992, §2.9]). Then the solution  $x^H$  is interpolated to the next finer level by an *interpolation* operator  $I_H^h$  and used to update the solution  $x^h$  there.

In the geometrical context underlying most PDEs, a hierarchy of coarser levels is easily constructed by discretizing the PDE onto grids with increasing grid spacing, i.e., onto fewer sampling points.

The propagation of the residual from fine to coarse and then of the solution back from coarse to fine is called a V-cycle (Fig. 1). It needs  $O(n)$  time, since the size of the levels decreases exponentially. For suitable  $I_H^h, I_h^H$  and  $A^H$  it reduces the error by a constant factor [Briggs, 1999].

<sup>2</sup>We follow the literature on multigrid methods in distinguishing different levels by superscript  $h$ . For the transition between two levels  $h$  denotes the finer and  $H$  the coarser level.

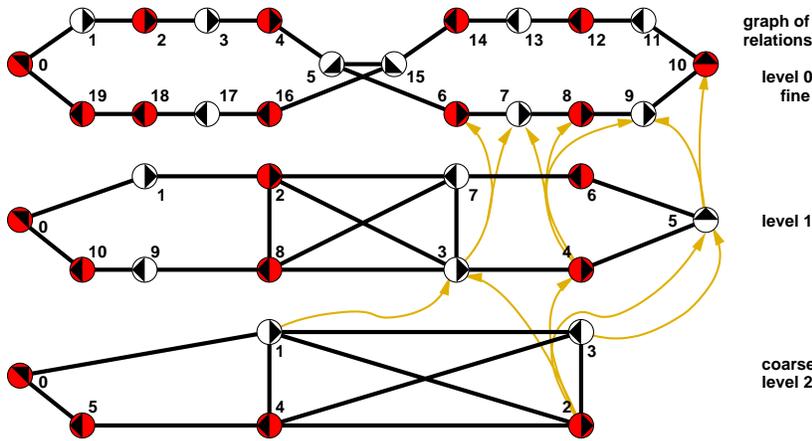


Figure 2: Example for a three level hierarchy: For each coarser level every even numbered frame is selected and the odd ones are interpolated from their even predecessor and successor. Coarse frames (black) are represented on the next coarser level, fine frames (white) are interpolated. Lines show the sparsity pattern of  $A^0$ ,  $A^1$  and  $A^2$ . The first level is identical to the graph of relations  $\mathcal{R}$ . Arrows show the interpolator  $I_1^0$  and  $I_2^1$  (only for frames 6...10)

### 3.2 Galerkin Multigrid

For PDEs,  $A^H$  can be naturally derived as the discretization onto a smaller set of sampling points.  $I_h^h$  and  $I_h^H$  are usually chosen as linear interpolation and weighted averaging respectively. If no “natural” choice for  $A^H$  and  $I_h^H$  is available, the Galerkin operator defines them purely algebraically for a given interpolator  $I_h^h$ . It is derived from the equivalent minimization problem (which on the finest level is just the original problem of minimizing  $\chi^2(x)$ ):

$$g(x) = x^{hT} A^h x^h - 2x^{hT} b \quad (8)$$

Since the coarse  $x^H$  corresponds to the fine  $I_H^h x^H$ , the coarse equation must minimize  $g(I_H^h x^H)$ :

$$0 = \frac{1}{2} \frac{\partial (g(I_H^h x^H))}{\partial x^H} \quad (9)$$

$$= \frac{1}{2} \frac{\partial (x^{HT} I_H^{hT} A^h I_H^h x^H - 2x^{HT} I_H^{hT} b)}{\partial x^H} \quad (10)$$

$$= \underbrace{I_H^{hT} A^h I_H^h}_{A^H} x^H - \underbrace{I_H^{hT} b}_{I_H^H} \quad (11)$$

So by using the Galerkin operator  $I_h^H = I_H^{hT}$ ,  $A^H = I_H^{hT} A^h I_H^h$ , the coarse equation minimizes  $g(x)$  over the range of the interpolator. Relaxation on any level thereby reduces  $g(x)$ , ensuring convergence to the unique solution for any  $I_H^h$ . For fast convergence, however, the choice of  $I_H^h$  is still crucial.

Another point to consider is that  $I_H^h$  has to be local in some sense, otherwise coarser matrices will become increasingly dense, taking more than  $O(n)$  time per iteration.

### 3.3 Algebraic Multigrid

There exist so-called algebraic multigrid approaches that define the interpolator in a purely algebraic form without any reference to an underlying geometry or PDE [Stüben, 1999]. In principle these approaches appear advantageous for a problem with an irregular geometry such as

SLAM. We implemented a variant of the so-called “direct” interpolation [Stüben, 1999, §4.2]. It interpolates a frame  $i$  so that the result satisfies  $(Ax)_i = b_i$  given all coarse  $x_j$  and given linear interpolates for the fine  $x_j, j \neq i$ . For our data, this approach led to unacceptably dense matrices (§5). We therefore replaced it by a problem specific interpolator described in §4.3.

## 4 Multilevel Relaxation

In this section we describe the Multilevel Relaxation algorithm proposed in this paper.

Unlike many PDEs, in SLAM the problem is not discretized onto a regular grid, so the question is how to define the hierarchy of coarser levels. The algorithm exploits the fact that the frames form a sequence, namely the robot’s trajectory, so selecting every second frame is a suitable way of generating a coarser level (Fig. 2). It uses a multilevel representation for equation (6) with a sparse matrix  $A$ . On this hierarchy it implements a Galerkin based V-cycle. The algorithm is incremental, updating  $\hat{x}$  for each new frame. Such an update involves three tasks: (i.) Extend  $A^h, b^h$  on all levels necessary to represent the new frame. (ii.) Update  $A^h, b^h$  and  $I_H^h$  based on the new relations. (iii.) Apply  $c$  V-cycles to update the ML estimate  $\hat{x}$ . The first two steps involve only few entries of  $A^h$  and  $b^h$  and take  $O(\log n)$ , the third step takes  $O(cn)$ .

### 4.1 Data Structure

The algorithm maintains the graph of relations  $\mathcal{R}$  in the usual way, with linked lists that allows efficient traversal of the set of edges incident upon a given node. Each relation  $r$  stores the corresponding Gaussian  $\mu^r, C^r$  and linearization point  $\check{a}^r, \check{b}^r$  (§2.1).

For the multilevel hierarchy, each level  $h$  contains the sparse equation matrix  $A^h$ , vector  $b^h$ , overall solution  $\hat{x}^h$ , residual solution  $x^h$  and the sparse interpolation matrix  $I_H^h$ .  $A^h$  is stored as a set of  $3 \times 3$  blocks  $\{(i, j, A_{ij}) | A_{ij} \neq 0\}$ . Blocks of a given row are linked for efficient traversal. This allows computation of  $(Ax)_i = A_{i\bullet} x$  and relaxation by equation (7) in  $O(1)$ .

Let $\mathcal{F}^h$ be the set of all frames involved in a new relation $r \in \mathcal{R}'$
Extend $\hat{x}^h, \check{x}^h$ with initial estimates for new frames
Define linearization points $\check{a}^r, \check{b}^r$ for new relations $r \in \mathcal{R}'$ (§2.1)
Recompute rows $\mathcal{F}^h$ of $A^h, b^h$ by (5) for finest level $h$
FOR Level $h$ from fine to coarse
Recompute rows $\mathcal{F}^h$ of $I_H^h$ by (19)
Let $\mathcal{F}^H = C(\mathcal{F}^h)$ be the set of affected coarse frames
Copy $\hat{x}_f^H$ from corresponding fine frame $\hat{x}_{F(f)}^h$ for all $f \in \mathcal{F}^H$
Recompute rows $\mathcal{F}^H$ of $A^H = I_H^{hT} A^h I_H^h$
Compute Cholesky decomposition $A^H = U^T U$ for coarsest level $H$

Figure 3: Update of multigrid hierarchy after adding new relations  $\mathcal{R}'$ .

The interpolator  $I_H^h$  is stored in an array, since each row contains at most two blocks. We define  $C(f)$ , the set of coarse frames from which  $f$  is interpolated, and  $F(f)$ , the fine frame corresponding to  $f$ , as:

$$C(f) = \begin{cases} \{\frac{f}{2}\} & f \text{ even} \\ \{\frac{f+1}{2}\} & f \text{ odd} \wedge \text{last} \\ \{\frac{f-1}{2}, \frac{f+1}{2}\} & \text{otherwise} \end{cases}, \quad (12)$$

$$C(\mathcal{F}) = \bigcup_{f \in \mathcal{F}} C(f), \quad F(f) = \begin{cases} 2f - 1 & \text{last} \\ 2f & \text{else} \end{cases} \quad (13)$$

## 4.2 Update

When new measurements arrive, a new frame is introduced into  $A^h, b^h$  and  $I_H^h$ , new relations are added and the equation is updated on each level. Only  $O(1)$  frames are involved, so the update is performed in  $O(1)$  per level and  $O(\log n)$  total.

Our approach is to always recompute a complete row of  $A^h, b^h$ , and  $I_H^h$ , keeping track of the changed rows from fine to coarse level (Fig. 3). Since each row contains only  $O(1)$  nonzero blocks, updating only the affected blocks is little faster and more complicated. For a new frame  $A^h, b^h$  and  $I_H^h$  are extended as necessary.

Let  $\mathcal{F}^h$  be the set of frames adjacent to a new relation. From the sparsity pattern in (5) it can be seen that only  $A_{ij}^h, b_i^h$  for  $i, j \in \mathcal{F}^h$  changes, so we recompute rows  $\mathcal{F}^h$ . For each  $f \in \mathcal{F}^h$  the blocks generated by all relations incident to  $f$  are added to  $A_{f\bullet}^h$  and  $b_f^h$ .

Next the interpolator  $I_H^h$  is updated. From the Galerkin principle we are free to choose the interpolator, so we update rows  $\mathcal{F}^h$ , to limit the resulting change in  $A^H$ . Equation (14) shows the structure of  $I_H^h$ :

$$I_H^h = \begin{pmatrix} I & & & \\ E_1^- & E_1^+ & & \\ & I & & \\ & E_3^- & E_3^+ & \\ & & I & \\ & & & \ddots \end{pmatrix}, \quad \begin{aligned} (I_H^h)_{ij} &\neq 0 \\ \Leftrightarrow j &\in C(i) \end{aligned} \quad (14)$$

Row  $(I_H^h)_{f\bullet}$  has blocks at columns  $C(f)$ , so updating rows  $\mathcal{F}$  changes columns  $C(\mathcal{F})$ . For a coarse  $f$ , row  $(I_H^h)_{f\bullet}$  is a single identity block  $I$ . For a fine  $f$ , it is two

blocks  $E_f^+$  and  $E_f^-$  defined by (19). The last step is to update  $A^H = I_H^{hT} A^h I_H^h$ , resulting in a change to columns  $C(\mathcal{F})$  of  $I_H^h$  and  $\mathcal{F}$  of  $A^h$ . A change in  $A_{ij}^h$  changes rows  $C(i)$  of  $A^H$ , so that rows  $C(\mathcal{F})$  are recomputed. When a level  $H$  has less than  $n_{\min}$  frames (32 in our experiments), the equation on that level is solved directly using a Cholesky decomposition  $A^H = U^T U$  computed during update.

## 4.3 Interpolation

There are several difficulties in devising a good interpolator in the context of SLAM: (i.) The interpolator must be based on the matrix  $A^h$  or estimate  $\hat{x}^h$ , not on the set of relations, since the latter is only available at the finest level. (ii.) It must be rotation invariant, since otherwise it creates apparent orientation information in the coarse equations, since for some orientations the interpolation fits better than for others. Since orientation is usually very uncertain [Frese and Hirzinger, 2001], this effect distorts the coarse solution. (iii.) This may even happen for rotation invariant interpolators due to linearization of the rotation. To see this, we consider the following relation and substitute into (3):

$$\mu = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, C = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \check{a} = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \check{b} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix} \quad (15)$$

$$\chi^2 \begin{pmatrix} a \\ b \end{pmatrix} = (a_x - b_x - 1)^2 + (a_y - b_y - b_\phi)^2 + (a_\phi - b_\phi)^2 \quad (16)$$

If, for instance, the rotation invariant interpolator  $b = a$  is chosen,  $\chi^2$  simplifies to  $1 + b_\phi^2$ , representing apparent absolute orientation information. To reduce this problem, we use the following geometric formula for interpolating a fine frame  $b$  from coarse frames  $a$  and  $c$ :

$$b = a + \alpha(c - a) + \beta(c - a)^\perp, \quad (17)$$

$$\alpha \in [0 \dots 1], \beta \in [-1 \dots +1], \quad (18)$$

$$\begin{pmatrix} b_x \\ b_y \\ b_\phi \end{pmatrix} = E_b^- a + E_b^+ c, \quad (19)$$

$$E_b^- = \begin{pmatrix} 1-\alpha & \beta & 0 \\ -\beta & 1-\alpha & 0 \\ 0 & 0 & \frac{1}{2} \end{pmatrix}, \quad E_b^+ = \begin{pmatrix} \alpha & -\beta & 0 \\ \beta & \alpha & 0 \\ 0 & 0 & \frac{1}{2} \end{pmatrix} \quad (20)$$

$\chi_0^2 = \chi^2(\hat{x}); \quad ctr = 0; \quad i = 0$	
WHILE $ctr < 3$	
Update $\hat{x}$ by 3 V-cycles, $\chi_i^2 = \chi^2(\hat{x}); \quad i = i + 1$	
$\gamma = 0.1 \frac{n}{m-n} \quad \alpha = \sqrt{\chi_i^2 / (\gamma \chi_0^2)} \quad \hat{\chi}_{\min}^2 = \chi_{i-1}^2 - \frac{\chi_{i-1}^2 - \chi_i^2}{1-\alpha}$	
IF	$\chi_i^2 \geq \chi_{i-1}^2 \vee (\chi_i^2 < (1 + \gamma) \hat{\chi}_{\min}^2)$
THEN	Update $\check{a}^r, \check{b}^r$ for all $r \in \mathcal{R}$ ; $ctr = ctr + 1$
ELSE	$ctr = 0$

Figure 4: Computation of the approximate ML estimate  $\hat{x}$ .

It defines the vector  $b - a$  as a linear combination of  $c - a$  and the orthogonal vector  $(c - a)^\perp$ . Therefore it is rotation invariant. The constants  $\alpha$  and  $\beta$  are chosen, so that  $E_b^- \hat{a} + E_b^+ \hat{c} = \hat{b}$ , but clipped to avoid extreme cases. Thereby the position of  $b$  relative to  $a$  and  $c$  closely matches the position used for linearization, and the above mentioned problems are reduced.

#### 4.4 Nonlinearity and Convergence

To obtain a consistent estimate incrementally, a single V-cycle for each new frame appears to suffice (§5), even when closing a loop. We update the linearization point  $\check{a}^r, \check{b}^r$  of a portion of the relations afterwards (5% in our experiments), so that the map can converge to the *non-linear* ML estimate while the robot continues moving. This is a great advantage over EKF based implementations, which do not allow changing of the linearization point after integration and can thus be subject to severe linearization errors [Frese and Hirzinger, 2001].

For an immediate ML estimate  $\hat{x}_{ML} = \arg \min_x \chi^2(x)$ , iteration with a termination criterion is performed (Fig. 4). The idea is to stop when the equation error  $\hat{x} - \hat{x}_{ML}$  is much smaller than the estimation error  $\hat{x}_{ML} - x_{\text{true}}$ . We estimate  $\hat{\chi}_{\min}^2 \approx \min_x \chi^2(x)$ , by assuming exponential convergence. The convergence factor  $\alpha$  is computed from the initial and last  $\chi^2$  values<sup>3</sup>.

It is well known that the expected minimum  $E(\min_x \chi^2(x))$  is  $3(m - n)$  and the expected  $E(\chi^2(x_{\text{true}}))$  value is  $3m$  [Press *et al.*, 1992, §15.1]. So  $\frac{n}{m-n} \hat{\chi}_{\min}^2$  is a rough estimate for  $\chi^2(x_{\text{true}}) - \min_x \chi^2(x)$ . When  $\chi^2(\hat{x}) < (1 + \gamma) \hat{\chi}_{\min}^2$  with  $\gamma = 0.1 \frac{n}{m-n}$  the linearization points  $\check{a}^r, \check{b}^r$  are updated, usually leading to further reduction of  $\chi^2(\hat{x})$ . If this happens three times in a row, iteration is stopped.

## 5 Results

We have evaluated the performance of the proposed algorithm on two well known datasets, one from the University of Freiburg [Gutmann and Konolige, 1999] and a single loop taken from the Carnegie Mellon Wean

Hall [Thrun *et al.*, 1998]. They are processed by the software package *ScanStudio*<sup>4</sup> which performs the scan-matching. The resulting graph of relations is passed to our implementation, computes the  $\chi^2$  function and uses either ‘Cholesky decomposition’ (CD)<sup>5</sup>, ‘Single level relaxation’ (SLR), ‘Multilevel relaxation’ (MLR) or ‘One MLR iteration’ (1MLR) for minimization. All four algorithms start with an initial estimate based on the first relation involving a frame. The last two methods ‘Incremental Multilevel Relaxation’ (IMLR) and ‘Incremental Single Level Relaxation’ (ISLR) apply a single MLR and SLR iteration for each new frame. Thereby they incrementally maintain an estimate as our algorithm would actually be used on a mobile robot. All experiments were conducted on a Pentium IV, 1.7 GHz using LINUX/gcc 2.95.3 (Fig. 5, 6).

For both datasets, MLR is much more efficient than CD and provides a better estimate. The latter point is true because CD solves the linearized problem, while all others perform nonlinear minimization. It is worth noting that linearization effects can be seen in the  $\chi^2$  value despite the small orientation error. SLR is faster than MLR on the Freiburg data, but much slower on the Wean Hall data. The reason therefore lies in the difference between the two datasets (Fig. 7a, c). The Wean Hall data is a long loop with a large global error. MLR is more efficient in reducing this type of error than SLR, which needs many more iterations. The error in the Freiburg data is mainly local, so both MLR and SLR need the same number of iterations.

There is an inconsistency in the lower right two rooms of the Freiburg CD estimate (Fig. 7a), which is also visible for the MLR, SLR and 1MLR estimates. The reason is that scans from the lower room and scans from the upper room overlap only slightly through the small doorway, so *ScanStudio* did not match any of them, and this inconsistency is not visible in the graph of relations.

IMLR and ISLR are much faster than CD, MLR and SLR, if an incremental estimate is desired. For the Freiburg data, both estimates are extremely good (Fig.

<sup>3</sup>The factor  $\gamma$  is a heuristic for the case that  $\chi_0^2$  is already very close to the minimum.

<sup>4</sup>We would like to thank Steffen Gutmann for the Freiburg data and the permission to utilize *ScanStudio*, and Sebastian Thrun for the Wean Hall data.

<sup>5</sup>a direct  $O(n^3)$  equation solver [Press *et al.*, 1992, §2.9]

	Freiburg			Wean Hall		
	iter.	time	$\chi^2$	iter.	time	$\chi^2$
Initial Estimate			16395061			1126227
Cholesky Decomposition (CD)		19.951 s	428397		1.268 s	6113
Single level relaxation (SLR)	12	0.437 s	431995	630	0.786 s	7122
Multilevel relaxation (MLR)	12	0.586 s	427178	12	0.059 s	5992
One MLR iteration (IMLR)	1	0.023 s	501273	1	0.003 s	40375
Exact Minimum			425639			5986
Incremental MLR (IMLR)	1	avg. 14.4 ms	426104	1	avg. 1.6 ms	6178
Incremental SLR (ISLR)	1	avg. 8.6 ms	425759	1	avg. 0.7 ms	91772
$n, m$	906	8081		346	932	
Blocks $\neq 0$ in $A^0, A^1, A^2$	15770	9824	4154	2054	848	414

Figure 5: Performance on Freiburg / Wean Hall data: CD, SLR, MLR, 1MLR all compute a batch estimate for the whole data set. IMLR and ISLR incrementally process each new frame (the average time per frame is given). The exact minimum was computed by iterating MLR to numerical convergence of the equation  $Ax = b$ .

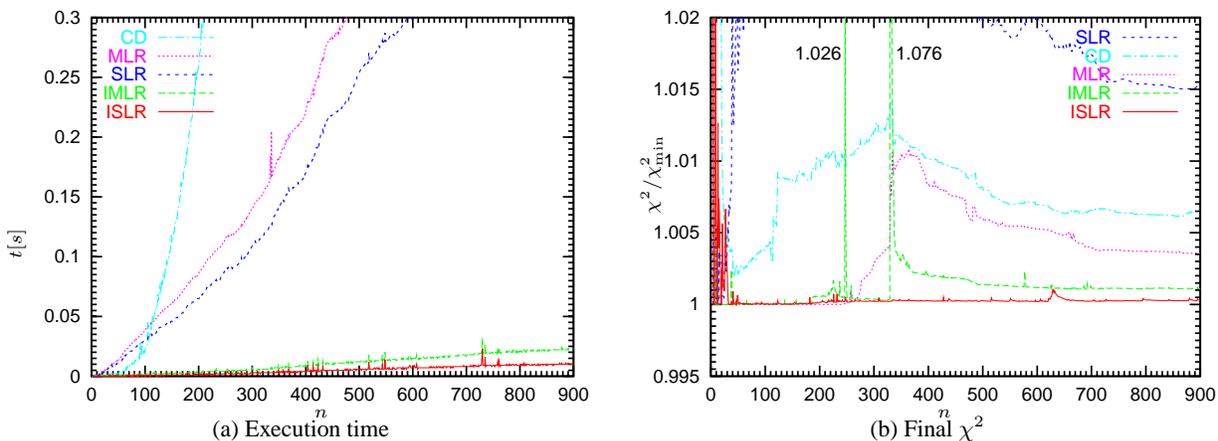


Figure 6: Performance on Freiburg data plotted over number of frames. Algorithm names are sorted from high to low values. (b) is scaled to show  $\chi^2$  values up to 2% above the minimum, which all correspond to excellent estimates.

6, 7b) and better than CD and MLR most of the time, with the exception of two outliers occurring after integrating two inconsistent relations. Surprisingly, the ISLR estimate is even better than the IMLR estimate, which has been found to be related to linearization effects. For the Wean Hall data, the CD and MLR estimates are initially better than the IMLR estimate (Fig. 6, 7d-e), which is in turn much better than the ISLR estimate. This is because both perform only a single iteration after closing the loop. Here the advantage of IMLR can be seen, since it closes the loop consistently (Fig. 7e), which is not achieved by ISLR.

The interpolator leads to sufficiently sparse matrices  $A^0, A^1, A^2$ , with each coarser level having 40% to 60% fewer nonzero blocks. When using a variant of *direct interpolation* [Stüben, 1999],  $A^1$  has 29723 and  $A^2$  has 25347 nonzero blocks, which is unacceptably dense.

The minimum  $\chi_{\min}^2$  is much larger (Freiburg:  $\times 20$ , Wean Hall:  $\times 3.4$ ) than the theoretically expected value

$3(m - n)$ . This shows that the scan matching covariance is overconfident [Bengtsson and Baerveldt, 2001] and stresses the importance of defining the termination criterion *relative* to  $\chi_{\min}^2$  in §4.4.

## 6 Conclusions and Future Work

This paper introduced a new SLAM algorithm, Multi-level Relaxation, which is suitable for incremental, on-line use on a mobile robot in  $O(n)$  time, including closing of large loops. This is possible because (i.) the algorithm makes an iterative refinement to the existing solution at each step, rather than re-solving the equation system from scratch, and (ii.) it exploits an important property of multigrid methods, namely that the residual error is geometrically smooth, i.e., it is distributed evenly over the whole map. In the case of closing a very large loop, as in the Wean Hall example presented, it can take several further iterations to converge to the maximum likelihood solution. However, the map is already geometri-

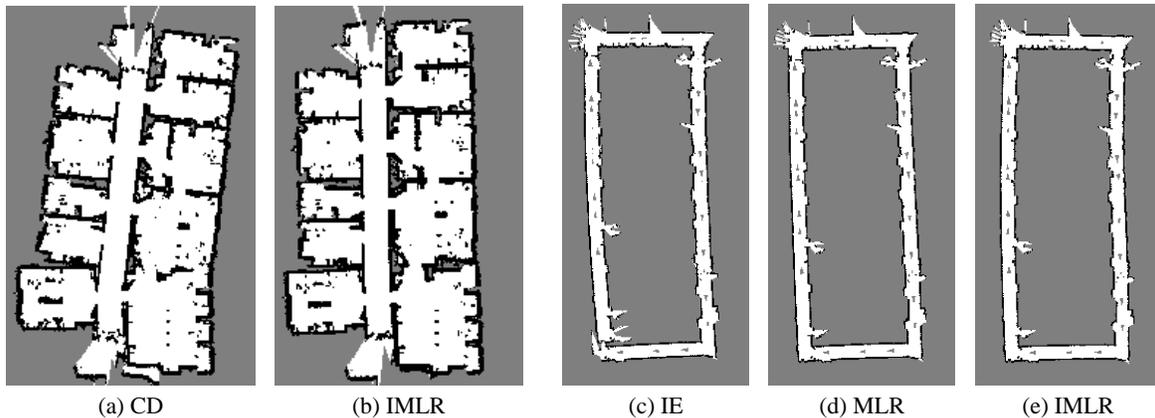


Figure 7: Computed maps for Freiburg ( $17m \times 26m$ , a-b) and Wean Hall ( $36m \times 74m$ , c-e).

cally *consistent* after a single iteration, that is, none of the measured relations are strongly violated in the estimated vector  $\hat{x}$ , and the map should be useful for navigation purposes. A further advantage of relaxation methods is that non-linearities can be handled by recomputing the linearization points if necessary. Remarkably, the result from a few iterations is already better than the exact solution of the linearized problem provided by Cholesky decomposition.

Future work will include embedding the new algorithm in a framework for handling both the continuous and discrete uncertainty in the SLAM problem. This would be achieved by multi-hypothesis tracking in the space of possible maps, where one hypothesis corresponds to one possible topology. In this paper, we have only used the algorithm to solve the linear equation system for a single topological hypothesis (i.e., we assumed no data association errors), but it should also provide a good core engine for multi-hypothesis SLAM, e.g., by tracking the best  $M$  topological hypotheses, due to its ability to close loops efficiently. While loops occur relatively rarely in most indoor environments, alternative topological interpretations of the same sensor data within a multi-hypothesis framework will often correspond to decisions on whether or not to close a loop – this is why an efficient equation solver is highly desirable for solving the SLAM problem in its most general form.

## References

- [Bengtsson and BaerVELdt, 2001] O. Bengtsson and A.-J. BaerVELdt. Localization in changing environments - estimation of covariance matrix for the IDC algorithm. In *Proc. IROS, Hawaii, USA*, pages 1931 – 1937, 2001.
- [Brandt, 1977] A. Brandt. Multi-level adaptive solutions to boundary-value problems. *Math. Comp*, 31:333–390, 1977.
- [Briggs, 1999] W.L. Briggs. A multigrid tutorial. Ninth Copper Mountain Conference On Multigrid Methods, April 1999. (<http://www.llnl.gov/CASC/people/henson/mgtut/ps/mgtut.pdf>).
- [Duckett *et al.*, 2002] Tom Duckett, Stephen Marsland, and Jonathan Shapiro. Fast, on-line learning of globally consistent maps. *Autonomous Robots*, 12(3):287–300, 2002.
- [Frese and Hirzinger, 2001] Udo Frese and Gerd Hirzinger. Simultaneous localization and mapping - a discussion. In *Proc. IJCAI Workshop on Reasoning with Uncertainty in Robotics, Seattle*, August 2001.
- [Guivant and NeBOT, 2001] J.E. Guivant and E.M. NeBOT. Optimization of the simultaneous localization and mapping algorithm for real-time implementation. *IEEE Trans. Robotics and Automation*, 17(3):242–257, 2001.
- [Gutmann and Konolige, 1999] J.S. Gutmann and K. Konolige. Incremental mapping of large cyclic environments. In *Proc. IEEE Int. Symposium on Computational Intelligence in Robotics and Automation - CIRA-99, Monterey, CA*, November 1999.
- [Lu and Milios, 1997] F. Lu and E. Milios. Globally consistent range scan alignment for environment mapping. In *Autonomous Robots*, volume 4, pages 333 – 349, 1997.
- [Montemerlo *et al.*, 2002] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. FastSLAM: A factored solution to the simultaneous localization and mapping problem. In *Proc. AAAI Nat. Conf. on Artificial Intelligence*, Edmonton, Canada, 2002.
- [Press *et al.*, 1992] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery. *Numerical Recipes, Second Edition*. Cambridge University Press, 1992. ISBN 0-521-43108-5.
- [Stüben, 1999] K. Stüben. Algebraic multigrid (AMG): An introduction with applications. Technical Report 70, GMD - Forschungszentrum Informationstechnik, November 1999.
- [Thrun *et al.*, 1998] S. Thrun, W. Burgard, and D. Fox. A probabilistic approach to concurrent mapping and localization for mobile robot. *Machine Learning*, 31(5):1 – 25, 1998.
- [Thrun *et al.*, 2002] S. Thrun, D. Koller, Z. Ghahramani, H. Durrant-Whyte, and Ng. A.Y. Simultaneous mapping and localization with sparse extended information filters. In *Proc. of the Fifth Int. Workshop on Algorithmic Foundations of Robotics*, Nice, France, 2002.