To appear in Proc. of the 6th Int. IFAC Symposium on Robot
Control (SYROCO'00), Vienna, Austria, September 2000.

BUILDING GLOBALLY CONSISTENT GRIDMAPS FROM TOPOLOGIES

Tom Duckett [*] **Alessandro Saffiotti** [*]

[*] *Center for Applied Autonomous Sensor Systems*
Department of Technology
University of Örebro
SE-70182 Örebro, Sweden
`http://www.aass.oru.se`
`{Tom.Duckett,Alessandro.Saffiotti}@aass.oru.se`

Abstract: This paper addresses the problem of recovering metric consistency in a global gridmap for mobile robot navigation. Gridmaps can only be updated consistently using exact estimates of the robot position, a requirement which is very hard to fulfil in real world environments because the same sensor data must be used for both map building and self-localisation. To overcome this problem, we use a hierarchy of robot maps which integrates topological and grid-based representations at different levels of abstraction. The consistency problem is solved at the topological level, by applying a relaxation technique to generate coordinates for the places in the robot's map. Consequently, the robot is able to recover a globally consistent gridmap without requiring accurate sensors or high computational costs. Experiments on a Nomad 200 robot in a large, real world environment are presented which demonstrate the efficacy of the approach.

Keywords: Mobile Robot Navigation, Map Building, Hybrid Maps, Occupancy Grids, Topological Maps, Relaxation Algorithm, Spatial Semantic Hierarchy.

1. INTRODUCTION

Maps are essential for mobile robot control in unstructured environments, being needed for self-localisation, path planning and human-robot interaction. A popular representation paradigm for robot maps is the occupancy grid model, see e.g., (Moravec and Elfes, 1985; Hughes and Murphy, 1992; Oriolo *et al.*, 1998). In this approach, the map consists of a matrix of cells, each containing some measure of the certainty that the corresponding area of the environment is occupied by an object.

If a robot is to function autonomously, it needs the ability to build its own maps. This requirement imposes some severe practical problems for a robot attempting to construct a global grid model in real-time:

- *Dependence on accurate position information.* Large-scale gridmaps can only be updated consistently using exact estimates of the robot position. Under realistic operating conditions it is often very difficult to maintain the required level of accuracy. For example, some systems require *a priori* position information from an external agent (Fabrizi *et al.*, 2000). Others depend critically on accurate sensing, e.g., using laser-range finders (Yamauchi *et al.*, 1998) and stereo vision (Thrun *et al.*, 1998*a*), to reduce positioning errors.
- *High computational cost.* When accurate global position information is not available, the same sensor data must be used both to build the map and to update the robot's position. Most current approaches do so by applying some optimisation technique over the space of possible maps, e.g., (Thrun *et al.*, 1998*c*; Thrun *et al.*, 1998*b*). These solutions tend to require large amounts of memory and processing power. For example, the technique proposed in (Thrun *et al.*, 1998*c*) requires up to two hours of computation to generate a gridmap with a spatial resolution of 1 meter in a large environment (90 × 90 meters).

An alternative paradigm is provided by topological maps, where the environment is represented as a graph of connected places. In this approach, the problem of self-localisation becomes that of *place recognition* (Kortenkamp and Weymouth, 1994), and the robot does not need to know its precise Cartesian coordinate for map building. The compactness of topological representations also means that computational costs are much lower than for gridmaps. However, these maps do not provide a detailed geometric interpretation of the target environment.

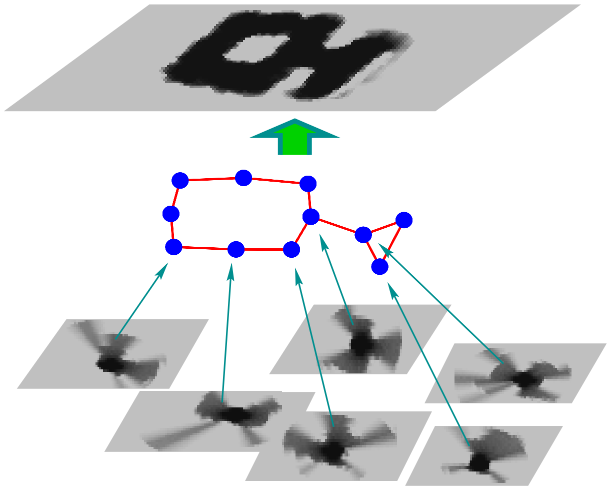

Fig. 1. Hierarchy of Maps.

In this paper, we integrate topological and grid-based representations for the purpose of constructing globally consistent metric models of large, real world environments. The core of the integration is the ability to obtain precise position information from a topological map without requiring accurate sensors and without incurring high computational costs. This information is then used to build a global gridmap. Using this approach, a sonar-equipped mobile robot is able to construct detailed models of large environments.

Our method relies on the combination of three existing mechanisms for robot map building:

- a topological map building strategy (Duckett and Nehmzow, 1999*a*),
- a relaxation technique for maintaining geometric consistency in a graph (Duckett *et al.*, 2000), and
- an off-line algorithm for constructing a global gridmap (Oriolo *et al.*, 1998).

The latter algorithm requires exact position information, which is obtained here by applying the relaxation technique to a self-acquired topological map.

Our method assumes four sources of perceptual information; (i) a place recognition system, (ii) a global orientation, (iii) local distance information, and (iv) range information used to build the gridmap. In this paper, we present experiments conducted on a Nomad 200 robot, in which we used the Bayesian self-localisation algorithm described in (Duckett and Nehmzow, 1999*b*) for (i), a compass for (ii), odometry for (iii), and sonar sensors for (iv). In the experiments, no *a priori* position information or map were provided to the robot.

2. METHOD

We integrate topological and grid-based representations at different levels of abstraction in a hierarchy of robot maps. The basic idea is summarised in Fig. 1.

At the lowest level, the environment is represented by a set of *local* grid models. Each of these grids constitutes a *perceptual signature* for one particular place in the environment; there is no requirement of consistency between the local grids. These local grids are used on-line by the robot for self-localisation.

At the intermediate level, the places are connected by a set of links to create a topological map. Each link is also labelled with metric information describing the relative distance and absolute angle between the two places it connects. Using this information, we then apply the relaxation technique described below to assign a geometrically consistent set of Cartesian coordinates to the places in the topological map.

Finally, the globally consistent metric information derived from the topological level is combined with recorded ultrasonic range data to generate a global gridmap.

These representations are manipulated by applying the following techniques. Detailed descriptions of these algorithms may be found in the papers available online (see References).

2.1 *Topological Map Building*

To obtain the topological map, an incremental map building strategy was applied (Duckett and Nehmzow, 1999*a*), in which the robot continuously tries to expand the territory which has already been charted. The basic idea is that the robot travels to the edge of the existing map, and then uses its range-finder sensors to detect more unexplored places. The new places are added to the map, then the process is repeated until the whole environment has been covered.

A particular feature of the approach is that an artificial neural network is trained to predict the presence of unexplored places in a given direction, fusing together information from the robot's range-finder sensors, see (Duckett and Nehmzow, 1999*a*) for details. The new "predicted" places are added to the map, then subsequent movement by the robot is used to verify whether the predicted places actually exist or not.

During on-line operation, the robot maintains a temporary local grid model in working memory corresponding to its most recent sensory perceptions; this is used for collision avoidance, place recognition (by matching with the stored place signatures) and initialising the perceptual signatures of new places.

2.2 *Relaxation Algorithm*

A major problem for robot map building is that odometry-based dead reckoning cannot be used for accurate position estimation because of cumulative drift errors. To overcome this problem, we applied the iterative relaxation algorithm described in (Duckett *et al.*, 2000) to assign geometrically consistent position information to the places in the topological map. In this algorithm, the coordinates of the places are treated as free variables, and the relaxation method finds a globally consistent set of coordinates using only the local metric relations between places.

In this approach, each link in the topological map can be modelled as a spring which connects two adjacent places i and j, where each link is labelled with the relative distance d_{ij} and absolute heading θ_{ij} between places i and j. Each "spring" reaches minimum energy when the relative displacement between the coordinates of i and j is equal to the vector $(d_{ij},\ \theta_{ij})$ measured by the robot. Thus, global consistency is maintained in the map by minimising the following energy function:

$$E = \sum_i \sum_j{}' (x_i - x_j + d_{ij}\cos\theta_{ij})^2 + (y_i - y_j + d_{ij}\sin\theta_{ij})^2 \,,$$

where $\sum_j'$ refers to the sum over the neighbours of a given node i.

The basic principle behind the relaxation algorithm can be explained as follows. The idea is to pick each node in turn, and then move it "to where its neighbours think it should be" — see (Duckett *et al.*, 2000) for full details. By repeated application of this rule, the coordinates in the map converge towards a global minimum in the energy function. Furthermore, it has been proven that the algorithm converges to the maximum likelihood solution.

2.3 *Gridmap Construction*

The sonar data recorded at each place, together with the coordinates of that place after relaxation, are used to build a global occupancy grid by standard techniques. In our experiments, we have used the technique proposed by (Oriolo *et al.*, 1998), which is based on fuzzy logic. We had two main motivations in choosing this technique: (1) it maintains distinct maps for the occupied and the empty space, thus allowing us to distinguish between unexplored cells and cells on which there are contradicting measurements; and (2) it produces fuzzy gridmaps that can be processed by the technique proposed in (Fabrizi and Saffiotti, 2000) to extract higher level information, which can be used to further expand the map hierarchy. However, it should be noted that the approach proposed in this paper to obtain globally consistent gridmaps can be applied to any occupancy grid construction technique, e.g., (Moravec and Elfes, 1985; Hughes and Murphy, 1992).

3. EXPERIMENTS

Fig. 2. The Nomad 200 mobile robot *Milou*.

We have tested the above method in experiments performed using a Nomad 200 robot equipped with a compass and a ring of 16 Polaroid sonar sensors (Fig. 2). The experiments were conducted in the indoor environment shown in Fig. 3, which is a relatively large office area of size 46 × 12 meters.

Fig. 4 shows the topological map acquired by the robot in one experiment. The picture shows the position of the places in global coordinates before and after relaxation. The derived global gridmap is shown in Fig. 5. The map has a resolution of 0.10 meters, and should be accurate enough for safe navigation and planning.

The entire process requires minimal computational resources. Acquisition of the topological map was done on-line. Relaxation was performed as part of the acquisition algorithm. One iteration of relaxation on the full map of 137 places required 20 msec. Since the map was relaxed every time a new place was added to it, only one step was needed each time. The gridmap was generated off-line in these experiments, taking 36 sec to process 6600 sonar readings on a grid of 350 × 450 cells. All times are relative to a 200 MHz Pentium II processor.

4. CONCLUSIONS

Building a global gridmap requires exact position information. In on-line map building systems, this information is usually obtained by correcting the robot's odometry, e.g., using a Kalman filter (Gelb, 1974). However, the Kalman filter is based on assumptions which can be very hard to fulfil under realistic operating conditions.

By contrast, we have presented an application of a relaxation technique for building global gridmaps which is based on an underlying topological representation of the environment. Topological maps have the advantage that they can represent much larger areas using the same computational resources, and have a much lower dependency on accurate positioning and accurate sensing for map building.

The work presented in this paper belongs to a growing family of techniques that integrate map representations at different levels of abstraction and granularity. In many of these, the space is represented as a patchwork of locally consistent metric spaces connected to form a global topological map, e.g., (Duckett and Nehmzow, 1999*b*; Gasós and Saffiotti, 1999; Kuipers, 2000; Simhon and Dudek, 1998; Zimmer, 2000). Our work extends this type of approach by exploiting the information in the topological map to recover global metric consistency. In this respect, the closest relative of our method is the expectation maximisation (EM) technique proposed in (Thrun *et al.*, 1998*b*). While EM-based mapping techniques have produced impressive results, they suffer from a high computational complexity; moreover, EM is not guaranteed to converge to a global optimum. By contrast, in the case of the method described in this paper: (i) relaxation always finds a global optimum, and (ii) computational cost is low.

This work is part of our effort to develop an integrated hierarchy of robot maps. The different levels of the hierarchy may comprise representations with different semantics, abstractions and granularity. A key advantage of this approach is that it allows us to apply individual techniques appropriate to a particular level of the hierarchy, and thus to integrate disparate techniques for mobile robot navigation. One next step will be to use the generated global gridmap to derive higher level information, e.g., by applying the techniques for extracting morphological and semantic information concerning the structure of the space introduced in (Fabrizi and Saffiotti, 2000). So far, we have concentrated on integrating the layers in the hierarchy from the "bottom up"; future work will also investigate techniques for enhancing the functionality of the lower level navigation algorithms using higher level information.

5. REFERENCES

Duckett, T. and U. Nehmzow (1999*a*). Exploration of unknown environments using a compass, topological map and neural network. In: *Proceedings of the 1999 IEEE International Symposium on Computational Intelligence in Robotics and Automation,*. Monterey, CA. Online at http://aass.oru.se/˜tdt/.

Duckett, T. and U. Nehmzow (1999*b*). Knowing your place in real world environments. In: *Proceedings of EUROBOT '99, 3rd European Workshop on Advanced Mobile Robots*. IEEE Computer Press. pp. 135–142. Online at http://aass.oru.se/˜tdt/.

Duckett, T., S. Marsland and J. Shapiro (2000). Learning globally consistent maps by relaxation. In: *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'2000)*. San Francisco, CA. Online at http://aass.oru.se/˜tdt/.

Fabrizi, E. and A. Saffiotti (2000). Extracting topology-based maps from gridmaps. In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*. San Francisco, CA. Online at http://aass.oru.se/˜asaffio/.

Fabrizi, E., G. Oriolo and G. Ulivi (2000). Accurate map building via fusion of laser and ultrasonic range measures. In: *Fuzzy Logic Techniques for Autonomous Vehicle Navigation* (D. Driankov and A. Saffiotti, Eds.). Chap. 11. LNCS. Springer. Berlin, DE. To appear.

Gasós, J. and A. Saffiotti (1999). Integrating fuzzy geometric maps and topological maps for robot navigation. In: *Proc. of the 3rd Int. Symposium on soft computing*.

Gelb, A. (1974). *Applied Optimal Estimation*. MIT Press. Cambridge, MA.

Hughes, K. and R. Murphy (1992). Ultrasonic robot localization using Dempster-Shafer theory. In: *Neural and Stochastic Methods in Image and Signal Processing*. SPIE Vol. 1766.

Kortenkamp, D. and T. Weymouth (1994). Topological mapping for mobile robots using a combination of sonar and vision sensing. In: *Proceedings of the Twelth National Conference on Artificial Intelligence (AAAI'94)*. Seattle, Washington. pp. 979–984.

Kuipers, B. (2000). The spatial semantic hierarcy. *Artificial Intelligence* **119**, 191–233.

Moravec, H. and A. Elfes (1985). High resolution maps from wide angle sonar. In: *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'85)*. IEEE Computer Society Press. St. Louis, Missouri. pp. 116–121.

Oriolo, G., G. Ulivi and M. Vendittelli (1998). Real-time map building and navigation for autonomous robots in unknown environments. *IEEE Transactions on Systems, Man and Cybernetics* **28**(3), 316–333.

Simhon, S. and G. Dudek (1998). Selecting targets for local reference frames. In: *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'98)*. pp. 2840–2845.

Thrun, S., A. Bücken, W. Burgard, D. Fox, T. Fröhlinghaus, D. Henning, T. Hofmann, M. Krell and T. Schmidt (1998*a*). Map learning and high-speed navigation in RHINO. In: *AI-based Mobile Robots: Case Studies of Successful Robot Systems* (D. Kortenkamp, R.P. Bonasso and R. Murphy, Eds.). MIT Press. Cambridge, MA.

Thrun, S., J-S. Gutmann, D. Fox, W. Burgard and B. Kuipers (1998*b*). Integrating topological and metric maps for mobile robot navigation: a statistical approach. In: *Proc. of the 15th AAAI Conference*. pp. 989–996.

Thrun, S., W. Burgard and D. Fox (1998*c*). A probabilistic approach to concurrent mapping and localization for mobile robots. *Autonomous Robots* **5**, 253–271. Joint issue with *Machine Learning* 31:29–53.

Yamauchi, B., A. Schultz and W. Adams (1998). Mobile robot exploration and map-building with continuous localisation. In: *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'98)*.

Zimmer, U. (2000). Embedding local metrical map patches in a globally consistent topological map. In: *Proceedings of Underwater Technologies 2000*. Tokyo, Japan.

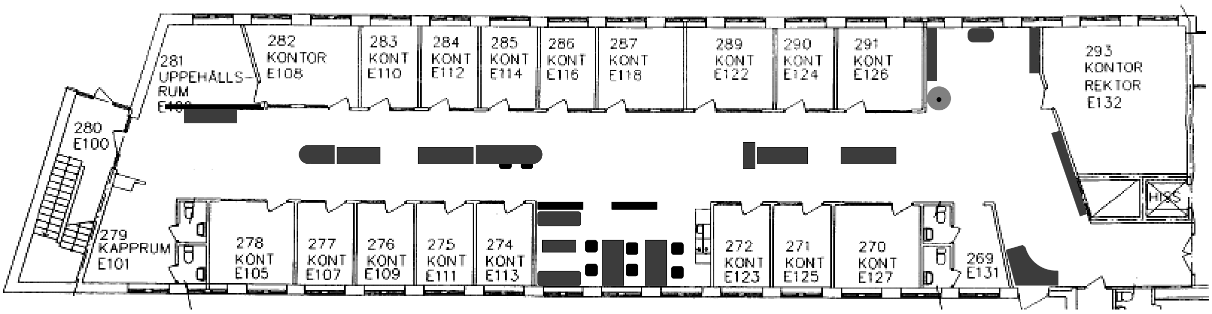

Fig. 3. The environment.

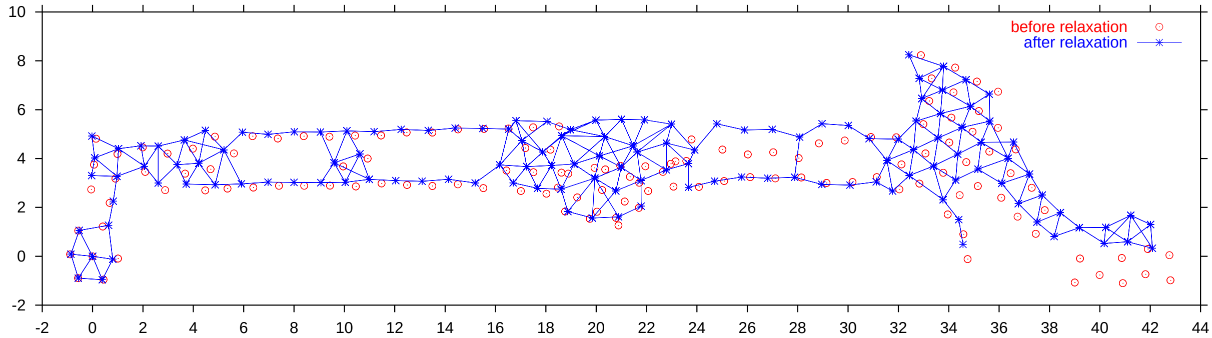

Fig. 4. The self-acquired topological map, showing coordinates before and after relaxation.

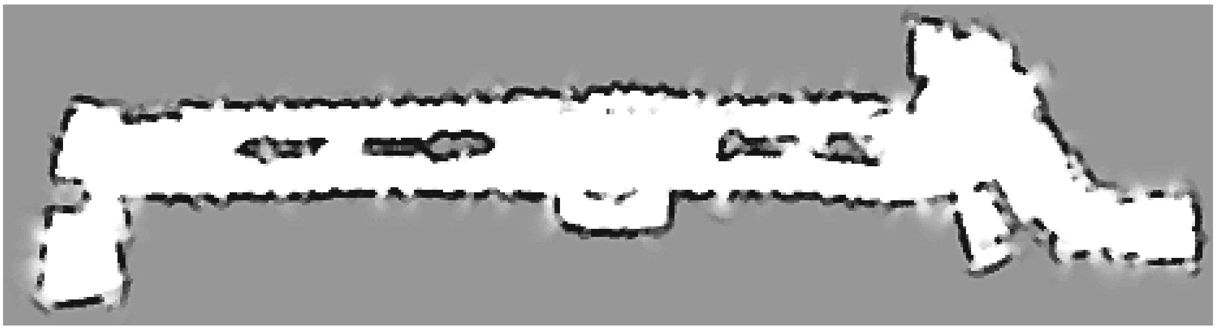

Fig. 5. The global gridmap.