

# Relaxation on a Mesh: a Formalism for Generalized Localization

**Andrew Howard, Maja J Matarić and Gaurav Sukhatme**

Robotics Research Labs, Computer Science Department,  
University of Southern California  
ahoward@usc.edu, mataric@usc.edu, gaurav@usc.edu

## Abstract

This paper considers two problems which at first sight appear to be quite distinct: localizing a robot in an unknown environment and calibrating an embedded sensor network. We show that both of these can be formulated as special cases of a *generalized localization* problem. In the standard localization problem, the aim is to determine the pose of some object (usually a mobile robot) relative to a global coordinate system. In our generalized version, the aim is to determine the pose of all elements in a network (both fixed and mobile) relative to an arbitrary global coordinate system. We have developed a physically inspired ‘mesh-based’ formalism for solving such problems. This paper outlines the formalism, and describes its application to the concrete tasks of multi-robot mapping and calibration of a distributed sensor network. The paper presents experimental results for both tasks obtained using a set of Pioneer mobile robots equipped with scanning laser range-finders.

## 1 Introduction

Over the last few years, we have witnessed the emergence and rapid maturation of a number of key embedded systems technologies, including reliable wireless communications and compact low-power micro-processors. These technologies are enabling the development of a wide range of novel *sensor/actuator networks*. We envisage networks containing a diverse collection of sensors and actuators, including devices that are fixed (surveillance cameras, motion sensors), devices that are carried (mobile phones, handheld computers) and devices that are mobile (robots). We assert that effective cooperation among such a diverse collection of elements will require some form of sensor fusion and/or joint planning and execution of actions. This will in turn require that the network elements have knowledge of their relative spatial configuration. The dynamic nature of this problem suggests that it can be formulated as a kind of *generalized localization* problem: in the standard localization problem, the aim is to determine the pose of some object (usually a mobile robot) relative to a global coordinate system; in our generalized version, the aim is to determine

the pose of all elements in the network (both fixed and mobile) relative to some arbitrary global coordinate system.

Consider, for a moment, a network containing only static elements. Assume that each element is equipped with either a *beacon* or *beacon sensor* such that the identity and pose of each beacon can be unambiguously determined by the beacon sensors. Each measurement made by a beacon sensor imposes a *constraint* on the relative pose of two network elements; given a set of such measurements, the generalized localization problem can be reduced to the task of finding a set of global poses (one for each network element) such that these constraints are satisfied. Our solution to this problem is based on a physical analogy with a mesh. Consider the system of rigid bodies connected by springs depicted in Figure 1. Imagine that each body corresponds to an element of the network and each spring corresponds to a constraint between elements. The springs are constructed in such a way that the energy stored in each spring is zero if the constraint is satisfied and greater than zero otherwise. By allowing this system to relax to its lowest energy configuration, we can determine the set of global poses that best satisfies the constraints.

Dynamic localization tasks can be solved in a similar manner. Consider a network whose elements are mobile and assume that these elements are equipped with either a beacon or a beacon detector. In addition, assume that each element is equipped with a motion sensor (such as odometry) that allows it to measure changes in pose. The solution to this localization problem can once again be found by creating a mesh of rigid bodies connected by springs. In this case, however, each network element is represented by a *series* of bodies, each of which describes the pose of that element at a particular time. Constraints arising from motion measurements are represented as springs between successive points in this series. By allowing the mesh to relax to its lowest energy configuration, we can determine the global pose of all network elements *at all times*. Note that this is particularly useful for map-building, since measurements are effectively propagated backwards in time; i.e., each new measurement will not only update the estimate of where elements are currently, it will update the estimate of where they were previously.

The mesh-based approach has a number of attractive features. It is general: localization, simultaneous localization and mapping, multiple robot simultaneous localization and mapping, and calibration of sensor/actuator networks can all be solved. It is robust: the result can be shown to be equivalent to a maximum likelihood estimator (that proof is, however, beyond the scope of this paper). It is scalable: the algorithm scales linearly with the number of elements being localized. It is simple: the core of the relaxation algorithm can be expressed in about 100 lines of C code. It is fast: the results described in this paper were generated in a few seconds using a Pentium III 450.

The mesh-based approach, as described in this paper, has two key limitations. First, it assumes that beacons have unique identities; the approach is thus unsuitable for problems involving natural landmarks, which inevitably have some degree of ambiguity associated with them. The approach is, however, very well suited to problems involving distributed sensor networks and/or teams of mobile robots; for such problems, we are free to ‘engineer’ both beacons and beacon sensors to provide unambiguous data. The second limitation of the mesh-based approach is that, for dynamic problems, the size of the mesh will grow linearly with time. While we believe this limitation can be mitigated though the intelligent deletion or merging of ‘older’ parts of the mesh, we have yet to implement any such mitigation strategies.

The remainder of this paper develops the mesh-based approach in somewhat more detail and shows how it can be applied to the concrete tasks of simultaneous localization and mapping (SLAM), multiple robot simultaneous localization and mapping (multi-SLAM), and calibration of a distributed sensor network. It also includes experimental results that validate the utility of the approach for each of these tasks.

## 2 Related work

Localization is an extremely well studied area in mobile robotics. Most approaches, however, can be classed as variants on two basic techniques: Kalman filters [1, 7, 9] and Bayesian/Markov localization [10, 11, 12]. While both techniques have been shown to support simultaneous localization and map building (sometimes with multiple robots), they have some unfortunate scaling characteristics: with Kalman filters, one must invert matrices of order  $n^2$ , where  $n$  is the number of elements to be localized; with Markovian systems, one must maintain probability distributions in  $3n$  dimensions. Consequently, the applicability of these approaches to localization tasks involving large numbers of dynamic elements is questionable. In contrast, the approach described in this paper scales linearly with  $n$ .

An alternative approach, described by Lu and Milios [8], seeks to obtain a globally consistent map by enforcing pairwise geometric relationships between individual laser

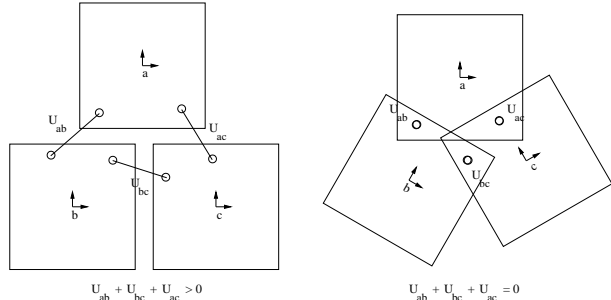


Figure 1: A simple mesh containing three rigid bodies and three springs. The mesh will relax to the configuration shown on the right.

range scans. The problem is framed as a maximum likelihood estimation problem in which the relationships between range scans are treated as random variables; the solution is computed directly through a sequence of (large) matrix operations.

In a similar vein, a number of authors have proposed localization methods based on relaxation of a mesh or truss [4, 2]. These share with Lu and Milios the notion of maintaining relative spatial relationships between locations or coordinate systems. The solution, however, is computed using an iterative technique, in which the overall system is allowed to ‘relax’ into the lowest energy state. The approach described in this paper falls into this class of solutions. Our formulation is, however, somewhat more general, allowing us to solve problems extending well beyond the canonical simultaneous localization and mapping example.

## 3 Formalism

### 3.1 Localization

Localization can be viewed as a coordinate transformation problem: every entity we wish to localize defines some local coordinate system (LCS) and every measurement we make defines a relationship *between* local coordinate systems. The aim of the localization process is to find the set of coordinate transformations that are consistent with these relationships.

Consider the case in which two different sensors measure the pose of the same object at the same time. While each sensor will measure this pose with respect to its own LCS, the fact that they are measurements of the *same* object means that this measurement can be used as a correspondence point. Let  $z'_a$  denote the pose of the object in LCS  $a$  and let  $z'_b$  denote the pose of the object in LCS  $b$ . We know that in some arbitrary global coordinate system, these two points must map onto the *same* point. Hence we can write:

$$\Gamma_a z'_a - \Gamma_b z'_b = 0 \quad (1)$$

where  $\Gamma$  is a coordinate transformation operator that maps points from the local to the global coordinate systems. The beacon and motion measurements described in the Introduction can both be cast into correspondence points of this form. For beacon measurements,  $z_a$  denotes the pose of the beacon as measured by the sensor and  $z_b$  denotes the pose of the beacon as measured *by the beacon*; the latter pose will be zero by definition. For motion measurements,  $z_a$  denotes the measured change in pose as the motion sensor moves from  $a$  to  $b$ , and  $z_b$  is again zero (by definition).

In the absence of uncertainty, Equation 1 can be solved directly to infer the coordinate transforms  $\Gamma$ . Real systems have uncertainty, however, so we must use some alternative technique. This is the motivation for the mesh-based approach: we represent each local coordinate system as a rigid body and each correspondence as a spring joining two points on these bodies (see Figure 1). By allowing this system to relax to its lowest energy configuration, we can infer the optimal set of coordinate transformations. Note that one can equally well think of this approach as a form of maximum likelihood estimation with gradient descent. We prefer the physical analogy, however, because of the intuitive insights it offers.

### 3.2 Relaxation on a mesh

Consider a pair of rigid bodies  $a$  and  $b$  that are joined by a spring. Let  $x_a$  denote the pose of body  $a$  in some arbitrary global coordinate system, and let  $x_b$  denote the pose of body  $b$  in the same global coordinate system. The spring joins the point  $z'_a$  on body  $a$  to the point  $z'_b$  on body  $b$ , where both  $z'_a$  and  $z'_b$  are defined in the local coordinate system of their respective bodies. The energy  $U_{ab}$  of this spring is given by:

$$U_{ab} = \frac{1}{2} k_{ab} (\Gamma_{x_a} z'_a - \Gamma_{x_b} z'_b)^2 \quad (2)$$

where  $\Gamma$  is a coordinate transform operator that maps points from local to global coordinate systems. The *spring constant*  $k_{ab}$  is typically set to  $1/\sigma^2$ , where  $\sigma$  is the uncertainty in the measurement represented by this spring. It is possible to formulate an expression for the spring energy that incorporates a full covariance matrix to represent measurement uncertainty. That topic is, however, beyond the scope of this paper.

Consider now a ‘mesh’ containing many rigid bodies, each of which is connected to other bodies in the mesh by one or more springs. The total energy of the mesh is given by:

$$U = \frac{1}{2} \sum_{ab} U_{ab} \quad (3)$$

where the sum is over all springs. Our aim is to find the set of poses  $\{x_a, x_b, \dots\}$  that minimizes this energy. A real mass-spring system will do this automatically: forces generated in each of the springs will push and pull the bodies

until the system settles into an equilibrium state in which all forces are zero. We can simulate this process using a simple relaxation algorithm:

- For each body  $a$ , compute the total force acting on that body:

$$F_a = \sum_b F_{ab} = - \sum_b \nabla_{x_a} U_{ab} \quad (4)$$

where  $F_{ab}$  is the force generated by the spring  $U_{ab}$  and  $\nabla_{x_a}$  denotes the gradient with respect to  $x_a$ . This gradient can be evaluated using the chain rule:

$$\nabla_{x_a} U_{ab} = J_{x_a} z_a \cdot \nabla_{z_a} U_{ab} \quad (5)$$

where  $z_a = \Gamma_{x_a} z'_a$  and  $J_{x_a} z_a$  denotes the Jacobi matrix of  $z_a$  with respect to  $x_a$ .

- For each mass  $a$ , update the pose  $x_a$  using the rule:

$$x_a \leftarrow x_a + \frac{1}{2} \Delta t^2 F_a \quad (6)$$

where  $\Delta t$  is an arbitrary time constant (which controls the rate of convergence).

This algorithm is repeated until the system reaches equilibrium. In theory, this occurs when the force on all nodes reaches zero; in practice, the algorithm is stopped when the magnitude of the total force on each node falls below some preset threshold.

Given that this algorithm is a form of gradient descent, there is some possibility that the system will converge to a local, rather than a global, minimum. While we have not found this to be a significant problem to date, the issue requires further investigation.

## 4 Experiments

### 4.1 Aim

We have conducted a series of experiments aimed at demonstrating both the validity and robustness of the mesh-based approach. Three experiments are described here: a simultaneous localization and mapping (SLAM) experiment, consisting of one robot equipped with a beacon detector and odometry in an environment containing fixed beacons; a multi-robot simultaneous localization and mapping experiment (multi-SLAM); and a calibration experiment, with one robot equipped with a beacon and odometry, in an environment containing fixed sensors. In all three experiments, we wish to validate that the maps generated by the relaxation algorithm are topologically consistent (in a global sense) and metrically accurate (in a local sense). Note that we have neither the means nor the motivation to measure the *global* metric accuracy of the map (global




Figure 2: A binary-coded laser beacon; the bright bars are retro-reflective tape.

metric accuracy is simply not required for most of the applications in which we are interested). We also wish to establish the *generality* of the formalism, which we have done by using the same program (applied to different data sets) in all three experiments.

Animations of the experiments, showing the map building and relaxation process, can be found at:

<http://robotics.usc.edu/~ahoward/projects.html>.

## 4.2 Apparatus and Methodology

The beacon detectors used in these experiments are standard SICK scanning laser range-finders. In addition to returning the range and bearing of objects, these devices can be programmed to output a 3-bit intensity value. Consequently, the laser can be configured to distinguish between objects of low and high reflectivity. We have constructed beacons that exploit this capability: the beacons are simple bar-codes in which strips of retro-reflective paper mark the ‘1’ bits and non-reflective paper mark the ‘0’ bits (see Figure 2). With a little post-processing of the laser signal, we are able to determine the range, bearing and orientation of these beacons from a distance of about 8m. The specific identity of each beacon can be determined from a distance of about 2m (the limiting factor being the the angular resolution of the laser).

The lasers are attached to fairly conventional ActivMedia Pioneer 2DX mobile robots running the *Player* robot server [3]. The server protocol allows for any number of remote clients to connect to a robot, or for a single remote client to connect to any number of robots. The experiments described here were performed in this latter configuration, with a single client collecting data from all of the robots. *Player* was developed at USC Robotics Research Labs and is freely available under the GNU Public License from

<http://robotics.usc.edu/player/>

## 4.3 Exp 1 – SLAM

For the first experiment, we attached six beacons at key locations in the corridors of an otherwise unmodified of-

fice building. The robot was ‘joysticked’ around the environment and the sensor data logged to a file. This file was later post-processed to generate the maps shown in the top row of Figure 3. These maps show the mesh overlaid on the raw laser-scan data. The circles in the mesh denote landmarks, the squares show the path of the robot and the links indicate measurements. Uncertainty in all measurements was assumed to be 10%.

The left hand figure shows the map before relaxation; this map effectively demonstrates the results using odometry only, with the odometric drift clearly in evidence. The right hand figure shows the map after relaxation; in this map, the odometric drift has been corrected by the beacon measurements. There are two features of this map that should be noted. Firstly, the relaxed map is highly rectilinear, despite the fact that the formalism described in this paper *does not* assume rectilinear environments. This result was unexpected, given the relatively sparse placement of beacons. Secondly, the relaxed map has corrected the systematic bias in the odometry, despite the fact that the formalism assumes that uncertainty in the measurements is unbiased. The quality of this result gives us some confidence in the robustness of the approach.

As a check of the metric accuracy of the relaxed map, we calculated the distance between topologically adjacent beacons and compared these distances to actual values obtained with a tape measure. The average difference in these distances was less than 7cm, corresponding to an error of less than 0.75%.

## 4.4 Exp 2 – Multi-SLAM

For the second experiment, three different robots were ‘joysticked’ through the same corridor environment used in Experiment 1. Data from all three robots was logged to a file, which was post-processed to generate the maps shown in the middle row of Figure 3. The left and right hand figures show the results before and after relaxation.

Since the robots have no means of detecting each other directly and start at different locations in the environment, they initially generate three separate maps. Eventually, however, the robots observe the same beacons, allowing them to merge their separate maps into a single shared representation. Once again, the relaxed map shows a high degree of rectilinearity and self consistency, despite the different biases present in each robot’s odometry.

The metric accuracy of the relaxed map was determined using the same procedure as in Experiment 1; the average error was found to be 13cm or less than 1.40%.

## 4.5 Exp 3 – Calibration

For the final experiment, we replaced each of the fixed beacons with a fixed beacon sensor and attached a beacon to a single mobile robot. We created, in effect, an ‘inverse’ SLAM experiment. Data from the fixed sensors and the

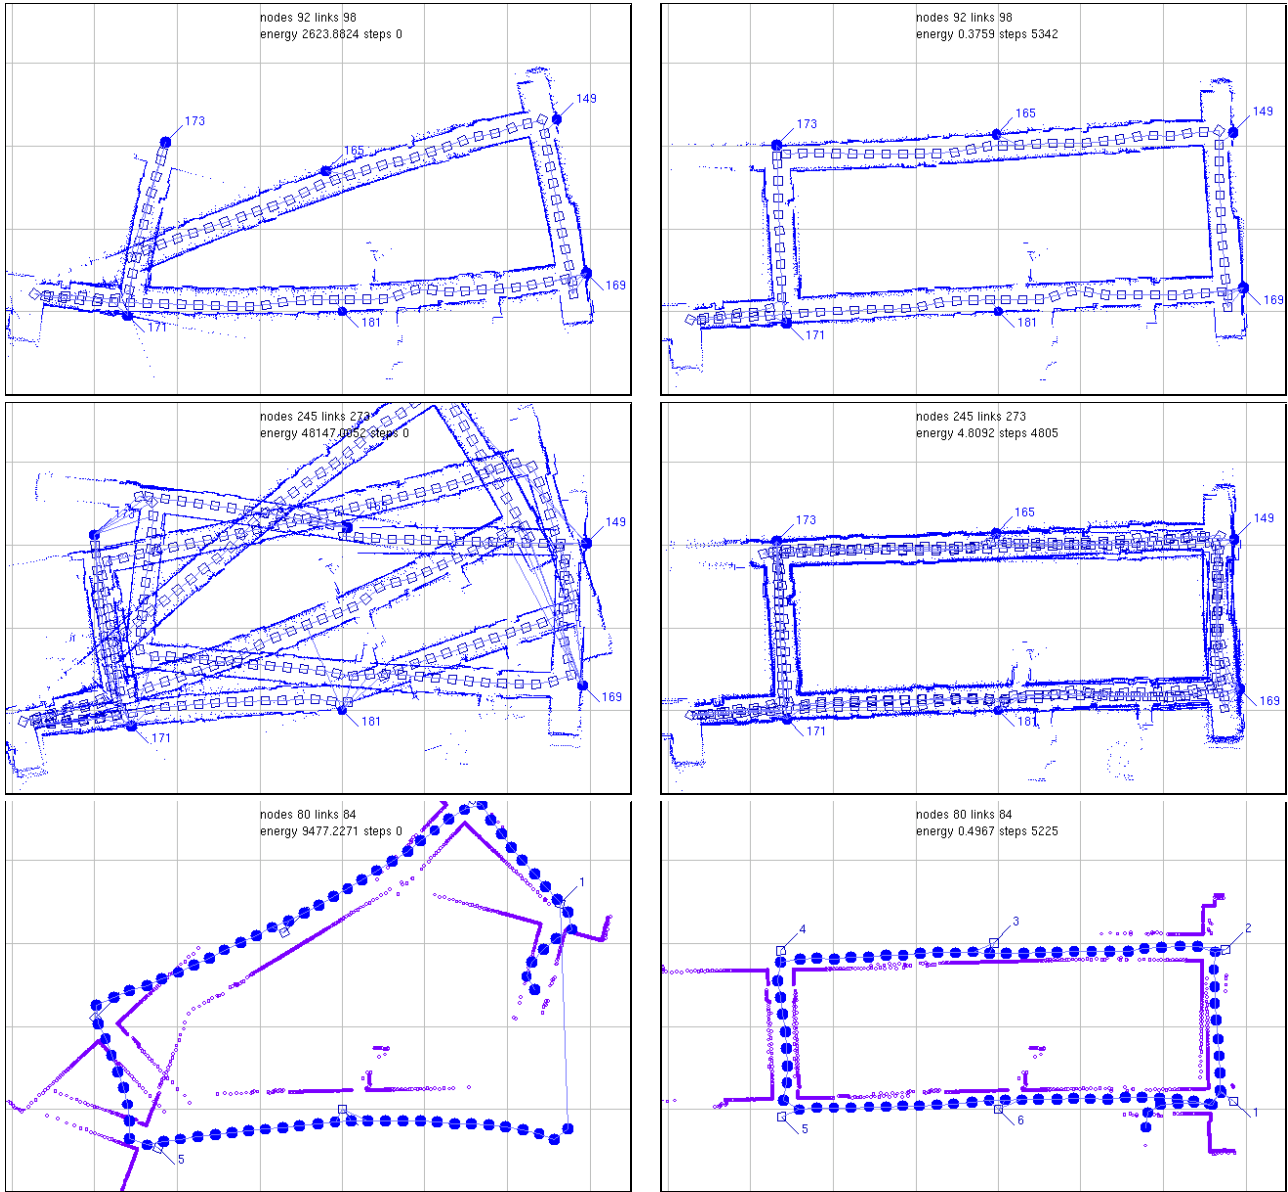


Figure 3: Experimental results. Top row: results for experiment 1 (SLAM), showing the map generated by a single robot. No information about the beacons was provided. The figures show the map before and after relaxation. Middle row: results for experiment 2 (multi-SLAM), showing the map generated by multiple robots. Bottom row: results for experiment 3 (calibration), showing the map generated by sensors emplaced in the environment. Note that this map shows live (not stored) data.

mobile robot was logged and post-processed to produce the maps shown in the bottom row of Figure 3.

Perhaps unsurprisingly, these maps look much like those produced in the previous two experiments. It is important to bear in mind, however, that while the first two maps display *stored* laser data, this map displays *live* laser data. One can readily use this data for applications such as monitoring and tracking. The quality of this map is remarkable when one considers that the sensors were not able to see

each other, nor were they able to detect the same beacon at the same time. The relationships between sensors are entirely indirect (via odometric measurements).

The metric accuracy of the relaxed map was measured, and was found to have an average error of less than 4cm or 0.31%. This is comparable to that obtained in the first two experiments.

## 5 Conclusion

While the experiments described in the previous section are far from exhaustive, they clearly show that the mesh-based formalism can solve generalized localization problems. Furthermore, they indicate that method is both general and robust.

There are a number of aspects of this formalism that are in need of further exploration. *Extended dynamics*: as noted in Section 3.2, the introduction of more complex dynamics into the relaxation process may result in faster convergence and/or allow the modeling of other aspects of the environment. *Global convergence*: in our experiments to date, the system has always relaxed to the global minimum. This does not, however, indicate that this must always be the case. Further work needs to be done to characterize the algorithm's convergence behavior. *Mixed beacon/sensor systems*: there are a number of potential applications that have yet to be tested experimentally. We would, for example, like to test the situation in which both sensors and beacons are on the robots, and the world is unmodified. Imagine a team of robots, each equipped with a beacon and a beacon detector, acting cooperatively to explore an environment [5, 6].

There are also many natural extensions to the work described here. *Anonymous beacons*: throughout this paper, we have assumed that beacons are uniquely identifiable. However, if we wish to apply the formalism to natural landmarks, we must allow for some degree of ambiguity (aliasing). We believe this can be done by adding a 'combinatorial layer' on top of the formalism described here. This layer would maintain a set of meshes, each of which corresponds to a different assignment of beacon identities. The total energy of each mesh after relaxation would be used to sort the meshes from the most probable to the least probable. *Multi-modal systems*: the approach needs to be extended to multi-model systems, so we can make concurrent use of multiple sensor types (such as laser range-finders and cameras). *Distributed mesh calculations*: ultimately, we would like to distribute the mesh maintenance and computation across multiple CPUs. The mesh can be divided into regions, each of which is assigned to a separate computer. The computers need only communicate information about the elements at the boundary of their regions. In this way, we hope to create a system that can scale to very large systems (with thousands or tens-of-thousands of elements).

## 6 Acknowledgements

This work is supported by the DARPA MARS Program grant DABT63-99-1-0015, ONR grant N000140110354, and ONR DURIP grant N00014-00-1-0638.

## References

- [1] S. Borthwick and H. Durrant-Whyte. Simultaneous localization and map building for autonomous guided vehicle. In *Proceedings of the IEEE/RSJ/GI International Conference on Intelligent Robots and Systems*, volume 2, pages 761–8, 1994.
- [2] T. Duckett, S. Marsland, and J. Shapiro. Learning globally consistent maps by relaxation. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 4, pages 3841–3846, 2000.
- [3] B. P. Gerkey, R. T. Vaughan, K. Støy, A. Howard, G. S. Sukhatme, and M. J. Mataric. Most valuable player: A robot device server for distributed control. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, Wailea, Hawaii, Oct. 2001.
- [4] M. Golfarelli, D. Maio, and S. Rizzi. Elastic correction of dead reckoning errors in map building. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 2, pages 905–911, 1998.
- [5] A. Howard and L. Kitchen. Cooperative localisation and mapping. In *International Conference on Field and Service Robotics (FSR99)*, pages 92–97, 1999.
- [6] R. Kurazume and S. Hirose. An experimental study of a cooperative positioning system. *Autonomous Robots*, 8(1):43–52, 2000.
- [7] J. J. Leonard and H. Durrant-Whyte. Simultaneous map building and localization for an autonomous mobile robot. In *Proceedings of the IEEE/RSJ International Workshop on Intelligent Robots and Systems*, volume 3, pages 1442–1447, 1991.
- [8] F. Lu and E. Milius. Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, 4:333–349, 1997.
- [9] S. I. Roumeliotis and G. A. Bekey. Collective localization: a distributed kalman filter approach. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 3, pages 2958–2965, 2000.
- [10] R. Simmons and S. Koenig. Probabilistic navigation in partially observable environments. In *Proceedings of International Joint Conference on Artificial Intelligence*, volume 2, pages 1080–1087, 1995.
- [11] S. Thrun, D. Fox, and W. Burgard. A probabilistic approach to concurrent mapping and localisation for mobile robots. *Machine Learning*, 31(5):29–55, 1998. Joint issue with *Autonomous Robots*.
- [12] B. Yamauchi, A. Shultz, and W. Adams. Mobile robot exploration and map-building with continuous localization. In *Proceedings of the 1998 IEEE/RSJ International Conference on Robotics and Automation*, volume 4, pages 3175–3720, 1998.