Simultaneous Localization and Environmental Mapping with a Sensor Network

Dimitri Marinakis, Neil MacMillan, River Allen, Sue Whitesides

Abstract— **In this paper, we present an algorithm for simultaneously refining a probability distribution function (PDF) for the *pose* of a sensor network (*i.e.* the locations of the sensors), and inferring the spatial variations of measured environmental parameters. Our approach iteratively refines a network pose PDF by assuming that environmental parameters vary smoothly. Both our physical experiments, which sensed wireless signal strength as the environmental variable, and our numerical simulations demonstrate that the approach has promise.**

I. INTRODUCTION

We present an Expectation Maximization (EM) algorithm for simultaneously learning both a probability distribution function (PDF) for the location of our sensors (*i.e.* the *pose* of the network) and also the spatial variations in parameters sensed in the surrounding environment. We assume that the parameters vary smoothly and that we have some prior knowledge of the pose. Our approach builds a model of how the environmental parameters measured vary over the region of interest and then uses this model to refine the estimate of where the sensors are located.

For example, suppose we are provided with a topographic map of the ocean bottom and know that a number of sensors have been deployed within the mapped region. Within this bounded region, the sensor positions are unknown, but each device can measure its depth (*e.g.*, with a pressure sensor). Using the topographic map, we could then build a prior probability distribution function (PDF) for the location of each sensor. Suppose each sensor additionally measures other properties in the environment such as salinity, temperature, current velocity, oxygen saturation, turbidity, light level, ambient magnetic field, etc. The question we ask is: can these other measurements be used to refine the PDF for the sensor locations? To answer the question we assume that the properties we sense are spatially correlated in some manner that we can model; *e.g.* the values vary smoothly.

Our algorithm works by simultaneously refining a probability distribution function (PDF) for the location of the sensors and a spatially dependent model for the parameter values. In the first iteration we bootstrap the process by using an initial PDF for each sensor location. Given this initial PDF, for each sensor we construct a multivariate field that provides a representation of the mean and uncertainty for each environmental property measured over the region of interest. This is done for a particular location by weighting the measurement of each sensor based on its likelihood of being at that location. The algorithm iterates over this process. The outcome is a refined PDF for the locations of the sensors as well as the inferred model of the environmental properties.

Computer Science Department, University of Victoria, BC, Canada dmarinak@gmail.com , neilrqm@gmail.com, riverallen@gmail.com, sue@uvic.ca

Note that there are a number of techniques that can be used to find a prior network pose PDF. Meger *et al.* [10] demonstrated that a network pose PDF can be obtained using the odometry data from a mobile robot through the use of Markov Chain Monte Carlo (MCMC). Also using MCMC-based techniques, Marinakis and Dudek [9] showed that a network pose PDF can be obtained using incomplete inter-sensor distance information, *e.g.* using time-of-flight methods. In this work we assume that such a prior has already been established and consider the problem of reducing the uncertainty associated with the pose PDF.

Main contribution: As far as we know, our technique of using smoothly varying parameters in the environment to aid localization in sensor networks is novel.

In the remainder of this paper, we first discuss some related work and then present a problem description and general formulation using the Expectation Maximization algorithm. We then give an example implementation of our approach and provide results both from simulations and from experiments that measured wireless signal strength as the environmental variable. We conclude with a discussion of our findings and future research directions.

II. RELATED WORK

Previous work by one of us used several MCMC instances to enable a sensor network to localize itself probabilistically [9]. A probabilistic method was used by Ihler *et al.* [8] to create a general particle filter using belief propagation, and by Blanco *et al.* [2] for large-scale Bayesian filtering. Expectation Maximization (EM) and its application to mixture densities have been described in [4], [11], [13]. Thrun [16] claims that EM is a consistent and orthogonal approach to localization and mapping, although it cannot generate maps incrementally. In [5] Detweiler *et al.* devise a simple method for mobile node localization using range or angle measurements. Their algorithm only requires knowledge of the positions of the anchor nodes and the maximum speeds of the mobile nodes.

Our experiments use cospatial *wifi* RSSI topography as environmental parameters. Low-power radio localization, which is usually based on mapping signal strength (RSSI) to distance [14], is unreliable due to obstacles, multipath fading, and noise [1]. In addition to RSSI, other methods for radio localization measure time difference of arrival or angle of arrival, but neither are in common use because they require special equipment, as noted in [6]. Graefenstein *et al.* [6] propose a system that maps RSSI to distance and direction measurements, which they conclude provides better localization compared to direction-only techniques. For the moment, we ignore the time-varying nature of the RSSI topography, but we note that Guha and Sarkar [7] address this issue in their model of signal strength in wireless networks. See, for example, Zhao and Guibas [21] for further details specific to terrestrial sensor network localization and related issues such as tracking.

With regard to underwater localization, Nawaz *et al.* [12] have designed a mobile sensor network for monitoring nuclear waste storage pools. They use a fleet of small submersible vehicles to explore and map the pools, and to detect and localize dangerous anomalies in the environment. In [20], Williams and Mahon use an underwater robot to perform SLAM using natural features in the coral of the Great Barrier Reef, and Walter *et al.* [19] propose a SLAM robot for inspecting ship hulls. Callmer *et al.* [3] describe a system for localizing using magnetometers, which sense landmarks with smooth magnetic topographies.

While our eventual interest is in underwater localization, our work can be applied to any environment with smoothly-varying environmental parameters. For example, Trincavelli *et al.* [18] use a robot to sense gas concentrations for pollution monitoring on land.

III. Problem Definition

We assume that we have N sensors deployed in the environment and that each of these sensors has a set of measurements of M distinct environmental properties. The observations can be described as a matrix O where each row o_i is a vector that gives the measurement of each property by sensor i. Additionally, we have a model that specifies how the M properties correlate spatially. Finally, we are given a prior PDF X^0 for the pose of the sensors.

The algorithm takes as input the observations O, the spatial correlation model, and the prior PDF X^0 for the sensor locations. The output is a refined PDF X for the location of the sensors and a model θ of how the M environmental properties vary over the region of interest.

IV. Simultaneous Localization and Environmental Mapping (SLEM)

We apply the Expectation Maximization (EM) algorithm [4] to the problem of simultaneously refining the pose of a sensor network and a model of how environmental parameters are spatially distributed. EM addresses the problem of fitting a model to data in cases where the solution cannot be easily determined analytically. The technique works by augmenting the existing observational data with unobserved, latent variables that can be used to incrementally improve the model estimate. The approach has been shown to converge to a set of model parameters that locally maximize the likelihood of the model, given the observations [4]. EM is a much used statistical technique and has been applied, for example, to key problems in robotics [15], [17].

We make the assumption that correlation between any two measurements in the environment is a function of the distance between the measurements. Specifically, in this work we assume that measurements taken from proximal locations are more strongly correlated than measurements take from distant locations.

A. EM applied to SLEM

To apply the EM algorithm to the SLEM problem for a sensor network, we iteratively refine an estimate for the network pose PDF X and our model θ of the environmental parameters. Here we give a generalized description of the problem formulation using the EM algorithm. In the next section we will provide an example of how the approach can be implemented.

We iterate over the following two steps:

1) *The E-Step:* define an auxiliary function Q that calculates the expected log likelihood of the complete data given the *last* estimate of our model, $\hat{\theta}$:

$$Q\big(\theta, \hat{\theta}\big) = E\Big[\log\big(p(O, X|\theta, X^0)p(\theta)\big)|O, \hat{\theta}\Big]$$

where O is our observed data, X^0 is our prior PDF for the network pose, and X, our latent or hidden data, is an estimate of the network pose PDF at the current EM iteration. A prior for the model is specified by $p(\theta)$.

2) *The M-Step:* update the estimate of θ with a value that maximizes our auxiliary function:

$$\theta = \operatorname*{argmax}_{\theta} Q\big(\theta, \hat{\theta}\big) \quad .$$

The auxiliary function Q can be expanded to the following:

$$Q\big(\theta, \hat{\theta}\big) = \int_x p(x|O, \hat{\theta}, X^0)\log\big(p(O, x|\theta)p(\theta)\big)dx \quad . \tag{1}$$

At each iteration of the EM algorithm we assign values to θ that maximize the log of the *expected* value of our observational data O and current estimate of the network pose PDF.

B. SLEM Implementation Details

The tractability of finding maximizing values for θ given the auxiliary function shown in Equation 1 depends on the representation of the PDF X and the specification of the prior over θ. By discretizing X and θ and assuming a uniform prior over θ, after some algebra, one can arrive at intuitive update equations (the problem then has similarities to typical EM algorithms such as multivariate Gaussian mixture density estimation). Specifying a prior over θ, however, that enforces the desired environmental smoothness assumption complicates the formulation. For the moment, we present an EM motivated algorithm for approximating a solution to the SLEM problem that has good performance in practice, but lacks a thorough derivation.

We represent the network pose PDF using occupancy grids. The environment is partitioned into a grid $G =$

$\{(k,l)\}$. The prior and current estimate of the network pose PDF X are then specified as a set of occupancy grids: $X = \{X_1, \ldots, X_N\}$ for each of the N sensors where $X_i = \{x^i_{kl}\}$ where x^i_{kl} gives the probability of sensor i being located within the grid cell $\{k,l\}$; *i.e.* $p(x_i = \{k,l\})$ where x_i denotes the grid location of sensor i.

Using Bayes law, and assuming that our prior for the network pose is independent of our model θ, we compute x^i_{kl} given the observation vector o_i and model θ as:

$$\begin{aligned} p(x_i = \{k,l\}|o_i, \theta) &= \frac{p(o_i|x_i = \{k,l\}, \theta)p(x_i = \{k,l\}|\theta)}{p(o_i|\theta)} \\ &\propto p(o_i|x_i = \{k,l\}, \theta)p(x_i = \{k,l\}) \end{aligned} \tag{2}$$

where $p(x_i = \{k,l\})$ is specified by our prior PDF X^0_i. The propotionality can then be converted into a probability by normalizing over all values in the grid X_i.

We now make the assumption that the location of each sensor is conditionally independent given the model of the environment θ. We model the environment θ as a multivariate field with a known mean and variance per each grid cell. More formally, $\theta = \{Z, P\}$ specifies a mean vector $Z = \{\mu_{kl}\}$ and covariance matrix $P = \{\Sigma_{kl}\}$ for the multivariate normal distribution associated with the cell (k,l).

We assume that a sensor measurement taken in a particular grid location is drawn from the density associated with that cell. Therefore, we compute the probability of observing the sensor vector o from sensor i given the model θ and location x_{kl} as:

$$p(o|x_i = \{k,l\}, \theta) = \mathcal{N}(o; \mu_{kl}, \Sigma_{kl}) \quad . \tag{3}$$

We can now compute the probability of our complete data, O and X, using Bayes law, and Equations 2, and 3:

$$\begin{aligned} p(O, X|\theta) &\propto p(X|\theta, O) \\ &\propto \prod_i^N \sum_{k,l \in G} x^i_{kl} p(o|x_i = \{k,l\}, \theta) \\ &\propto \prod_i^N \sum_{k,l \in G} x^i_{kl} N(o_i; \mu_{kl}, \Sigma_{kl}) \quad . \end{aligned} \tag{4}$$

In order to enforce the smoothness assumption, we specify that the maximum likelihood value for a field component takes proximal measurements into consideration. Given a PDF X and set of observations O, we compute a *weight factor* w^i_{kl} that gives the effect of the observation from device i on the grid location (k,l) as:

$$w^i_{kl} = \sum_{m,n \in G} x^i_{mn} \exp \frac{d(kl, mn)^2}{2L^2} \tag{5}$$

where $d(kl, mn)$ gives the Euclidean distance between the centres of grid cells (k,l) and (m,n) and L is a *smoothness parameter* that reflects how parameters in the environmental are assumed to vary as a function of distance. [1] The update equations for the parameters of θ given the current PDF estimate X and the observations O are then as follows:

$$\mu_{kl} = \frac{\sum_i^N w^i_{kl} o_i}{\sum_i^N w^i_{kl}} \tag{6}$$

$$\Sigma_{kl} = \frac{\sum_i^N w^i_{kl}(o_i - \mu_{kl})(o_i - \mu_{kl})^t}{\sum_i^N w^i_{kl}} \quad . \tag{7}$$

An EM iteration now consists of updating our PDF given the current estimate of our environmental model θ using Equations 2, and 3 and then computing a new estimate of the environmental model using Equations 5, 6, and 7. Algorithm 1 gives the details.

Algorithm 1 Simultaneous Localization and Environmental Mapping (SLEM)

```
INPUT: O = (o_1 ... o_N) { sensor observations}
INPUT: L { smoothness parameter }
INPUT: X^0  prior network pose PDF
INPUT: G  grid
θ^0 ⇐ uninformed
for  t = 1 to t = T  do
  for  i = 1 to i = N  do
    for  (k,l) ∈ G  do
      update x^i_kl) { Equations 2 and 3 }
    end for
    X^t_i ⇐ {x^i_kl},  (k,l) ∈ G
  end for
  for  (k,l) ∈ G  do
    for  i = 1 to i = N  do
      w^i_kl ⇐ f(X^t_i, L) { Equation 5 }
    end for
  end for
  for  (k,l) ∈ G  do
    update μ_kl and Σ_kl { Equations 6 and 7 }
  end for
  θ^t ⇐ {μ_kl, Σ_kl},  (k,l) ∈ G
end for
```

V. Simulation Results

To evaluate the SLEM algorithm we simulated a number of problem instances. In each problem instance, we created a smooth multivariate surface and discretized it into a grid. Each variate of a simulated surface was created by convolving a field of uniformly distributed values between 0 and 1 with a Gaussian. We then selected sensor locations uniformly at random on this surface. A measurement from each sensor was simulated by adding zero mean noise to the value of the multivariate surface at the grid location of the sensor. A prior X^0 for the network pose PDF was obtained by assuming that one of the surface parameters was known. Given this single variate surface and the sensor model, a PDF was obtained using Equation 2. Similarly, a 'best case' PDF was generated using the entire multi-variate surface for comparison purposes.

[1] For the moment we apply the same smoothness factor to each environmental parameter, but L could be parameter dependent .

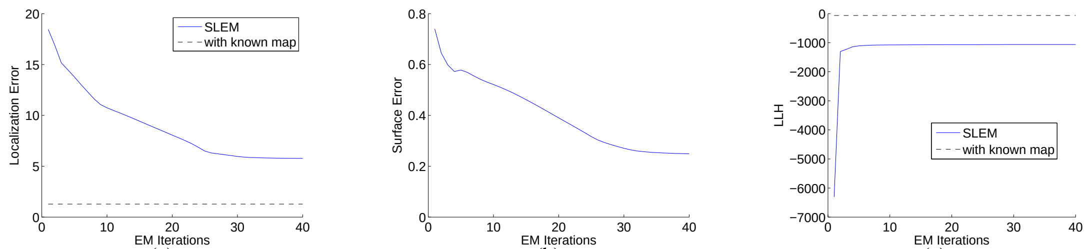

(a) (b) (c)

Fig. 2. Example of a problem instance with 100 sensors deployed on a 40x40 grid: (a) Localization Error; (b) Surface Error; (c) Log Likelihood. Same problem instance as shown in Figure 1.

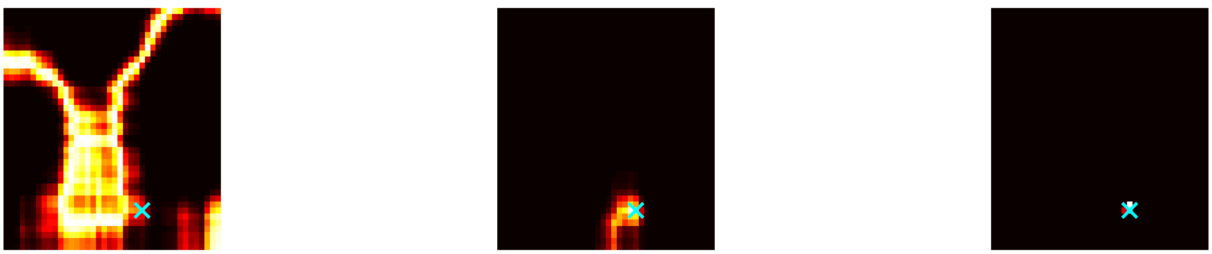

(a) (b) (c)

Fig. 3. Comparison of the PDF for a single sensor: (a) prior; (b) found by SLEM; (c) using the known multivariate surface. Lighter colours indicate higher probability; the blue X give the true location of the sensor. Same problem instance as shown in Figure 1.

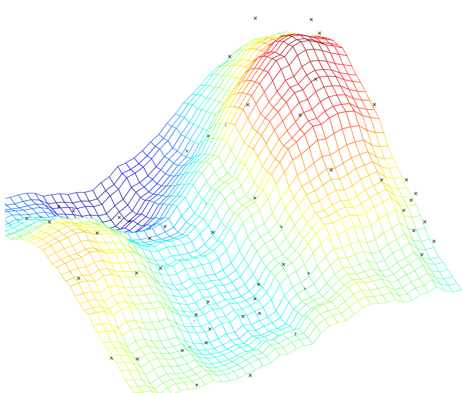

Fig. 1. Single variate surface with sensor locations overlaid as black crosses. The surface does not perfectly interpolate the sensor values due to simulated error measurements.

The SLEM algorithm was then given the prior network pose PDF X^0, the smoothness parameter L, and the measurements O. The output of the algorithm was a refined PDF X and an estimate of the multivariate surface obtained from the μ values of the inferred field θ.

A metric for the *localization error* in a network pose PDF given the true sensor locations was calculated as follows:

$$E_X = \frac{1}{N} \sum_{i}^{N} \left(\sum_{m,n \in G} x^i_{mn} d(mn, \lambda^i)^2 \right)^{1/2} \quad (8)$$

where x^i_{mn} gives the probability of sensor i being located in the cell (m, n) according to the PDF, λ^i gives the true grid location (k, l) of sensor i, and $d(mn, lk)$ gives the Euclidean distance between the centres of cells (m, n) and (l, k).

A metric for the estimated *surface error* based on θ and the true surface θ^* was calculated based on average squared error as follows:

$$E_\theta = \frac{1}{|G|} \sum_{m,n \in G} ||\mu_{mn} - \mu^*_{mn}|| \quad (9)$$

where μ_{mn} and μ^*_{mn} give the value at the grid cell (m, n) for the field θ and surface θ^*, respectively.

Our simulation results suggest that the SLEM approach shows promise as a technique for reducing network pose uncertainty. In one problem instance, 100 sensors were deployed in a region that was discretized into a 40x40 grid (see Figure 1). We simulated $M = 10$ different environmental parameters. Each variate of the surface was created by smoothing white noise with a Gaussian of standard deviation 8 grid units. Figure 2 shows the localization error, surface error, and likelihood as a function of each algorithm iteration for this problem instance. For reference, a network pose PDF was also obtained using the true multivariate surface. Figure 2(a) and (c) show the localization error and likelihood, respectively, for the PDF obtained using the true surface. Figure 3 shows the PDF for a single sensor.

The algorithm requires a suitable smoothness estimate. Figure 4 shows how the value selected for the smoothness

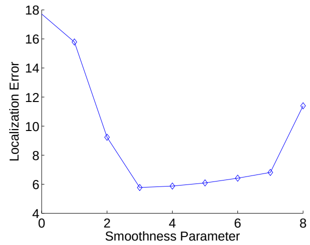

Fig. 4. Localization error as a function of the smoothness parameter L. Same problem instance as shown in Figure 1.

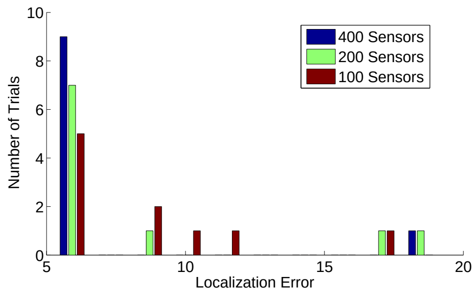

Fig. 5. Plot of final localization error for 10 different trials of problem instances on a 40x40 grid for different numbers of sensors.

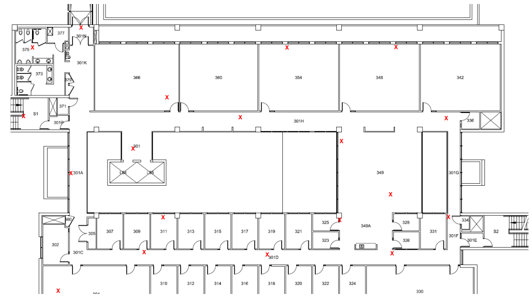

Fig. 6. Floor plan of office floor where data was collected with collection points marked with red crosses; region is 55 meters by 36 meters.

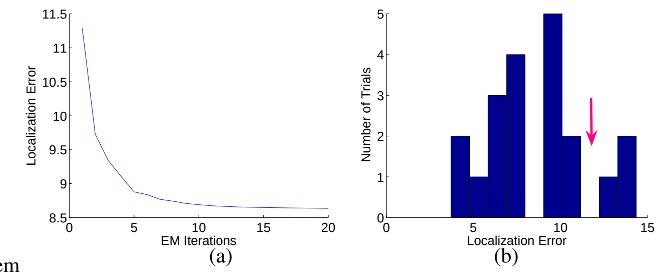

Fig. 7. Experimental localization error: (a) average location error as a function of algorithm EM iteration; (b) histogram of the location error per sensor, red arrow shows error based on prior PDF.

parameter L affected the localization error for this problem instance. Good results were obtained over a reasonably wide range of values.

It can be observed from Figure 2(c) that the solution found by the EM algorithm was considerably *lower* in likelihood than the solution obtained using the true surface, giving strong evidence that our current SLEM implementation can converge to a local optimum as opposed to the global optimum.

Trials over many problem instances reveal that although the algorithm generally improves the localization error, the results can be variable (Figure 5). When the density of the sensor deployment is increased, the consistency of the results improves. For example, half of ten trials resulted in a good localization error (near 5 units) for a 100 sensor problem type. This final error corresponds to approximately a factor of 3 reduction in uncertainty. When the trials were rerun using the same multivariate surface and problem setup, but with 400 sensors instead of 100, 9 out of 10 trials also resulted in a factor of 3 reduction in uncertainty.

Running our SLEM algorithm in un-optimized Matlab on a 100 sensor problem with a 40x40 grid took about 20 minutes using a single 2.4GHz CPU core.

VI. Experimental Results

We conducted a preliminary test of our approach using indoor *wifi* data. We measured the signal strength of a number of 802.11 access points at various locations throughout an indoor environment (Figure 6) over a period of roughly 2 hours. At each location, measurements were taken at 10Hz for a 30 second period and the median value was recorded. We assumed a prior PDF by centring a truncated Gaussian with a standard deviation of 8 meters around each measurement.

The SLEM algorithm was then used to refine the pose estimate based on the collected *wifi* data and provided prior PDF. Nine ($M = 9$) distinct access points that could be observed at each measurement point were selected as inputs to the algorithm. The region was discretized into a grid with a resolution of 1 meter. A value of $L = 4$ meters was selected for the smoothness parameter. (L values of 2 through 5 meters gave roughly comparable results.) Figure 7(a) shows how the average localization error was decreased by the algorithm from 11.3 meters to an average of 8.6 meters. There was, however, considerable variation in the localization error per sensor (Figure 7(b)). Of the 20 measurement locations, 17 had their location error reduced, while for 3 locations, the error was actually increased (Figure 8). Reducing the number of measurements used by the algorithm had the effect of lowering the localization accuracy, suggesting that denser measurements might have further improved the overall performance. There was also evidence that the algorithm was vulnerable to 'edge effects', as the devices that were localized least well were near the edges of the region.

VII. Discussion and Future Work

We have presented a framework for reducing positional uncertainty in sensor network pose estimation by assuming the parameters in the environment vary smoothly. We have

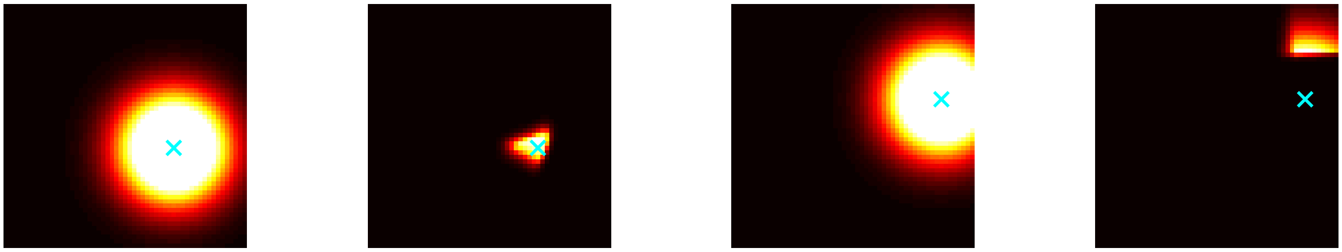

(a) (b) (c) (d)

Fig. 8. PDF for the best and worst located sensors using *wifi* RSSI: (a) prior PDF of best localized measurement; (b) final PDF of best localized measurement (3.7 meters); (c) prior PDF of worst localized measurement; (d) final PDF of worst localized measurement (14.3 meters).

presented a specific implementation using occupancy grids and a model for correlation in the environment that is based on an exponential function of distance. Our numerical simulations suggest that the approach can be effective at reducing uncertainty when observed environmental parameters vary smoothly. Experiments conducted using measured wireless signal strength verify the merit of the approach in a real environment.

One limitation of the approach thus far is that it may converge to a local maximum as opposed to a global maximum. Future work will look at applying other optimization techniques to this problem although even within the EM framework it remains to specify a suitable prior for the environmental model and cleanly derive the resulting update functions. [2] Additionally, the method should be extended to other representations of network pose PDFs since occupancy grids cannot be scaled to large environments. Further experimental validation and analysis is also required in order to help evaluate the merit of the approach.

Longer term research directions include the following. First, we plan to extend the correlation model to incorporate correlations that are not proximity-based, but rather based on the similarity of environment type. Second, we plan to investigate clustering similar environment types, the distribution of which could be evaluated given a prior. Finally, we plan to investigate the use of temporal variations in the sensor observations as additional environmental features.

Acknowledgments

We thank Mantis Cheng for technical assistance and lab space, and we acknowledge the Natural Sciences and Engineering Research Council of Canada for their funding.

References

[1] A. Awad, T. Frunzke, and F. Dressler. Adaptive distance estimation and localization in WSN using RSSI measures. 2007.

[2] J. Blanco, J. Fernandez-Madrigal, and J. Gonzalez. Toward a unified bayesian approach to hybrid metric-topological SLAM. *IEEE Transactions on Robotics*, 24(2):259–270, 2008.

[3] J. Callmer, M. Skoglund, and F. Gustafsson. Silent localization of underwater sensors using magnetometers. *EURASIP J. on Advances in Signal Processing*, 2010:1, 2010.

[4] A. Dempster, N. Laird, and D. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society*, 39:1–38, 1977.

[5] C. Detweiler, J. Leonard, D. Rus, and S. Teller. Passive mobile robot localization within a fixed beacon field. *Algorithmic Foundation of Robotics VII*, pages 425–440, 2008.

[6] J. Graefenstein, A. Albert, and P. Biber. Radiation pattern correlation for mobile robot localization in low power wireless networks. In *Proc. of the 2009 IEEE International Conference on Robotics and Automation*, pages 1260–1265, 2009.

[7] R. Guha and S. Sarkar. Characterizing temporal SNR variation in 802.11 networks. *IEEE Transactions on Vehicular Technology*, 57(4):2002–2013, 2008.

[8] A. Ihler and D. McAllester. Particle belief propagation. *AI & Statistics: JMLR W&CP*, 5:256–263, 2009.

[9] D. Marinakis and G. Dudek. Probabilistic self-localization for sensor networks. In *Proc. AAAI National Conference on Artificial Intelligence*, pages 976–981, Boston, Massachusetts, July 2006.

[10] D. Meger, D. Marinakis, I. Rekleitis, and G. Dudek. Inferring a probability distribution function for the pose of a sensor network using a mobile robot. In *Proc. of ICRA*, Kobe, Japan, May 2009.

[11] T. Moon. The expectation-maximization algorithm. *IEEE Signal Processing Magazine*, 13(6):47–60, 1996.

[12] S. Nawaz, M. Hussain, S. Watson, N. Trigoni, and P. Green. An underwater robotic network for monitoring nuclear waste storage pools. *Sensor Systems and Software*, pages 236–255, 2010.

[13] R. Redner and H. Walker. Mixture densities, maximum likelihood and the EM algorithm. *SIAM Review*, 26(2):195–239, 1984.

[14] A. Savvides, C. Han, and M. Strivastava. Dynamic fine-grained localization in ad-hoc networks of sensors. In *Proc. of the 7th Annual International Conference on Mobile Computing and Networking*, page 179. ACM, 2001.

[15] H. Shatkay and L. P. Kaelbling. Learning topological maps with weak local odometric information. In *International Joint Conference on Artificial Intelligence*, pages 920–929, San Mateo, CA, 1997.

[16] S. Thrun. Robotic mapping: A survey. *Exploring Artificial Intelligence in the New Millennium*, pages 1–35, 2002.

[17] S. Thrun, D. Fox, and W. Burgard. A probabilistic approach to concurrent mapping and localization for mobile robots. *Machine Learning and Autonomous Robots (joint issue)*, 1998.

[18] M. Trincavelli, M. Reggente, S. Coradeschi, A. Loutfi, H. Ishida, and A. Lilienthal. Towards environmental monitoring with mobile robots. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, 2008*, pages 2210–2215, 2008.

[19] M. Walter, F. Hover, and J. Leonard. SLAM for ship hull inspection using exactly sparse extended information filters. In *Proc. IEEE International Conference on Robotics and Automation*, pages 1463–1470, 2008.

[20] S. Williams and I. Mahon. Simultaneous localisation and mapping on the great barrier reef. In *Proc. IEEE International Conference on Robotics and Automation*, pages 1771–1776, 2004.

[21] F. Zhao and L. Guibas. *Wireless sensor networks: an information processing approach*. Morgan Kaufmann Pub., 2004.

[2] $p(\theta) = \prod_{ij} \sum_{kl \in G} e^{d(kl,ij)^2/2L^2} \mathcal{N}(\mu_{ij}; \mu_{kl}, \Sigma_{kl})$ might be an appropriate prior.