



This is a repository copy of *PipeSLAM: Simultaneous Localisation and Mapping in Feature Sparse Water Pipes using the Rao-Blackwellised Particle Filter*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/116928/>

Version: Accepted Version

Proceedings Paper:

Ma, K., Schirru, M., Zahraee, A.H. et al. (5 more authors) (2017) PipeSLAM: Simultaneous Localisation and Mapping in Feature Sparse Water Pipes using the Rao-Blackwellised Particle Filter. In: 2017 IEEE International Conference on Advanced Intelligent Mechatronics (AIM). IEEE International Conference on Advanced Intelligent Mechatronics (AIM), 2017, 3 - 7/7/2017, Munich, Germany. IEEE .

<https://doi.org/10.1109/AIM.2017.8014224>

© 2017 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other users, including reprinting/ republishing this material for advertising or promotional purposes, creating new collective works for resale or redistribution to servers or lists, or reuse of any copyrighted components of this work in other works.

Reuse

Unless indicated otherwise, fulltext items are protected by copyright with all rights reserved. The copyright exception in section 29 of the Copyright, Designs and Patents Act 1988 allows the making of a single copy solely for the purpose of non-commercial research or private study within the limits of fair dealing. The publisher or other rights-holder may allow further reproduction and re-use of this version - refer to the White Rose Research Online record for this item. Where records identify the publisher as the copyright holder, users can verify any specific terms of use on the publisher's website.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

PipeSLAM: Simultaneous Localisation and Mapping in Feature Sparse Water Pipes using the Rao-Blackwellised Particle Filter*

Ke Ma, Michele Schirru, Ali Hassan Zahraee, Rob Dwyer-Joyce, Joby Boxall, Tony J. Dodd, Richard Collins and Sean R. Anderson

Abstract—Water, a valuable resource, is usually distributed through urban environments by buried pipes. These pipes are difficult to access for inspection, maintenance and repair. This makes in-pipe robots an appealing technology for inspecting water pipes and localising damage prior to repair from above ground. Accurate localisation of damage is of critical importance because of the costs associated with excavating roads, disrupting traffic and disrupting the water supply. The problem is that pipes tend to be relatively featureless making robot localisation a challenging problem. In this paper we propose a novel simultaneous localisation and mapping (SLAM) algorithm for metal water pipes. The approach we take is to excite pipe vibration with a hydrophone (sound induced vibration), which leads to a map of pipe vibration amplitude over space. We then develop a SLAM algorithm that makes use of this new type of map, where the estimation method is based on the Rao-Blackwellised particle filter (RBPF), termed PipeSLAM. The approach is also suited to SLAM in plastic water pipes using a similar type of map derived from ultrasonic sensing. We successfully demonstrate the feasibility of the approach using a combination of experimental and simulation data.

I. INTRODUCTION

The maintenance of water infrastructure assets is of critical importance because water is such a highly valued resource. Water distribution pipes are usually buried and so are difficult to access. Therefore, robots have great potential for water pipe inspection. However, whilst a number of robot designs have been proposed for water pipes, along with sensors for inspection [1]–[3], the simultaneous localisation and mapping (SLAM) problem [4], for water pipes, has been much less widely studied. SLAM in water pipes is essential because damage must be accurately located to effectively target repair from above ground, minimising costs and disruption to both water supply and street users. Potentially, new ‘keyhole’ excavation techniques could be used for repair, greatly reducing disruption at street level and speeding repair, if the localisation of damage was sufficiently accurate. The main barrier to developing accurate SLAM techniques for water pipes is that the in-pipe environment is relatively

featureless. Therefore, the aim of this paper is to develop SLAM for feature sparse water pipes, increasing the features available for navigation.

Mapping and localisation has been developed for water pipes using SLAM techniques with cameras and inertial measurement units (IMUs) [5], [6]. A limitation of these methods is that there are few visual features available in the pipe. To alleviate this problem in plastic pipes, we have previously proposed the use of ultrasonics for measuring the soil profile through the pipe wall, which gives more features for navigation [7]. In effect, the depth observations from ultrasonic sensing form a continuous spatial map of the ground depth outside the pipe wall, which can be used with a terrain-based technique for localisation [8]. In that work, we assumed the map of terrain depth was known, in order to demonstrate the feasibility of localisation only, and in related work we also construct the map non-simultaneously [9]. This leaves a key gap for real-world problems: updating the map and simultaneously performing the localisation. In this paper we address the full SLAM problem for the first time in this type of scenario, and demonstrate the approach on metal water pipes.

To construct the map, a hydrophone is used to excite and measure metal pipe vibration, which is transformed to a one-dimensional function over space of pipe vibration amplitude. This provides a map analogous to those used in terrain-based localisation in aerospace [8] and SLAM in underwater robotics [10]. A terrain-type SLAM algorithm for water pipes, termed PipeSLAM, is developed here based on the Rao-Blackwellised particle filter (RBPF) [11]. Specifically, the map is regarded as a continuous function over space that is represented using a weighted basis function decomposition. We use a state-space model to represent robot pose and map parameter evolution. The corresponding state vector is well suited to estimation via the RBPF, because the map function is linear-in-the-parameters, and so can be estimated via the Kalman filter, and particles can be used to represent robot pose, solving the localisation problem.

The approach to mapping is demonstrated experimentally on a section of metal pipe, using a lab based setup. The PipeSLAM algorithm is evaluated using the experimental data as a ground truth map, with robot movement simulated over this map for offline analysis. The results of the PipeSLAM algorithm are benchmarked against a dead reckoning solution, and demonstrate that the RBPF algorithm can be used to more successfully solve the mapping and localisation problem using the terrain-type map.

*This work was supported by the EPSRC UK grant no. EP/K021699/1, Assessing the Underworld.

K. Ma, A. Hassan Zahraee, T. J. Dodd and S. R. Anderson are with the Department of Automatic Control and Systems Engineering, University of Sheffield, Sheffield, UK (phone +44 114 2225608; email: ke.ma@sheffield.ac.uk, a.hassanzahraee@sheffield.ac.uk, t.j.dodd@sheffield.ac.uk, s.anderson@sheffield.ac.uk).

J. Boxall and R. Collins are with the Department of Civil and Structural Engineering, University of Sheffield, Sheffield, UK (e-mail: j.b.boxall@sheffield.ac.uk, r.p.collins@sheffield.ac.uk).

M. Schirru, R. Dwyer-Joyce are with the Department of Mechanical Engineering, University of Sheffield, Sheffield, UK (e-mail: m.schirru@sheffield.ac.uk, r.dwyer-joyce@sheffield.ac.uk).

II. METHODS

In this section we develop the SLAM problem, estimation algorithm and describe the experimental details of mapping the metal pipe using hydrophone induced vibration.

A. Problem Statement

In this work we focus on the problem of robot navigation in relatively small pipes, with diameter of about 3 inches (a current prototype robot design is shown in Fig. 1). These small pipes are commonly used for water distribution in urban environments. The small diameter means that movement within the pipe is restricted, consisting of back and forth movement only. For entry into these small pipes, fire hydrants can potentially be used, which in the UK, Europe and the USA are spaced approximately every 100 metres [12, Chapter 14, Table 14.2].

We model the dynamics of the pipe robot using a state-space model, with state dynamics

$$p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_{k-1}) \Leftrightarrow \mathbf{x}_k = g(\mathbf{x}_{k-1}, \mathbf{u}_{k-1}) + \boldsymbol{\omega}_k \quad (1)$$

where k is the sample index, $\mathbf{x}_k \in \mathbb{R}^{n_x}$ is the robot pose (including location), $g(\cdot)$ is the state transition function, $\mathbf{u}_k \in \mathbb{R}^{n_u}$ is the input, $\boldsymbol{\omega}_k \sim N(0, \Sigma_w)$ is the state noise. In this paper we simplify the state vector \mathbf{x}_k to contain just the spatial location of the robot, x , along the pipe, and the input \mathbf{u}_k is obtained from a processed motor encoder reading defining a dead reckoning estimate of distance travelled, hence $g(\mathbf{x}_{k-1}, \mathbf{u}_{k-1}) = A_x \mathbf{x}_{k-1} + B_x \mathbf{u}_{k-1}$, where $A_x = 1$, $B_x = 1$. The measurement model is

$$p(\mathbf{y}_k | \mathbf{x}_k) \Leftrightarrow \mathbf{y}_k = h(\mathbf{x}_k) + \boldsymbol{\nu}_k \quad (2)$$

where the observation $\mathbf{y}_k \in \mathbb{R}^{n_y}$ is obtained from the hydrophone signal, where vibration amplitude is averaged over some frequency range, and the measurement function $h(\cdot)$ is a one-dimensional map, which transforms from the spatial location of the robot \mathbf{x}_k to mean pipe vibration amplitude; also $\boldsymbol{\nu}_k \sim N(0, \Sigma_v)$ is the measurement noise.

To represent the unknown map, $h(x)$, where x is a spatial location, we use a basis function decomposition,

$$h(x) = \sum_{j=1}^M \theta_j \phi_j(x) \quad (3)$$


where M is the number of basis functions, $\theta_j \in \mathbb{R}$ is the corresponding weight of the j th basis function, and $\phi_j(x)$ is a basis function. Note that this representation of $h(x)$ is linear-in-the-parameters θ_j . Here we use radial basis functions for $\phi(x)$, specifically the squared exponential function,

$$\phi_j(x) = \exp\left(-\frac{(x - c_j)^2}{2\sigma^2}\right) \quad (4)$$

where c_j is the centre of the basis function (a spatial location along the pipe) and σ is the width of a basis function, which we assume for simplicity here are the same across all basis functions. We also assume that the centres can be evenly spaced based on some prior knowledge of the likely spatial frequency of variation in pipe vibration amplitude, although



(a)



(b)

Fig. 1. Prototype of a water pipe inspection robot developed for this project. The robot is designed to operate in pipes with an approximate 3 inch diameter. The robot has a tether for power and communications, as well as to simplify recovery in the event of a robot failure (required by the water utilities). In order to move around bends the robot is comprised of small, separate modules that are linked by flexible steel springs. The core module units, of dimension 70 mm in length by 29 mm in diameter, are designed to contain sensors and processing units. The core units have flexible arms, which improve stability by bracing against the pipe wall. (a) Robot in 3 inch clear plastic pipe. (b) Zoomed view of one robot module.

it would be possible to use an adaptive technique to place basis functions [13]. Hence, in this case, the parameter vector $\boldsymbol{\theta}$ is the only unknown describing the map, where

$$\boldsymbol{\theta} = (\theta_1, \theta_2, \dots, \theta_M)^T \quad (5)$$

We define the probabilistic SLAM problem here in the usual way, as computing the joint distribution of robot location and map over all samples times [4],

$$p(\mathbf{x}_k, \boldsymbol{\theta} | \mathcal{Y}_k, \mathcal{U}_k, \mathbf{x}_0) \quad (6)$$

where $\mathcal{Y}_k = \{\mathbf{y}_1, \dots, \mathbf{y}_k\}$ is the set of observations, $\mathcal{U}_k = \{\mathbf{u}_1, \dots, \mathbf{u}_k\}$ is the set of all inputs and \mathbf{x}_0 is the initial location of the robot.

B. Estimation via a Rao-Blackwellised Particle Filter

To solve the problem of estimating the joint distribution of robot location and map parameters, $p(\mathbf{x}_k, \boldsymbol{\theta} | \mathcal{Y}_k, \mathcal{U}_k, \mathbf{x}_0)$, we initially define an augmented state-space model, with linear state dynamics

$$\begin{bmatrix} \mathbf{x}_{k+1} \\ \boldsymbol{\theta}_{k+1} \end{bmatrix} = \begin{bmatrix} A_x & 0 \\ 0 & I_M \end{bmatrix} \begin{bmatrix} \mathbf{x}_k \\ \boldsymbol{\theta}_k \end{bmatrix} + \begin{bmatrix} B_x \\ 0 \end{bmatrix} \mathbf{u}_k + \begin{bmatrix} 1 & 0 \\ 0 & I_M \end{bmatrix} \begin{bmatrix} \boldsymbol{\omega}_k \\ \boldsymbol{\eta}_k \end{bmatrix} \quad (7)$$

where I_M is the identity matrix of dimension M and $\boldsymbol{\eta}_k \sim N(0, \Sigma_\eta)$ is a noise term accounting for uncertainty in the map parameter evolution. The measurement model is

$$\mathbf{y}_k = \Phi(\mathbf{x}_k) \boldsymbol{\theta}_k + \boldsymbol{\nu}_k \quad (8)$$

where

$$\Phi(\mathbf{x}_k) = (\phi_1(\mathbf{x}_k), \dots, \phi_M(\mathbf{x}_k)) \quad (9)$$

The problem of estimating the augmented state vector is nonlinear, and so, for example an extended Kalman filter might be used to obtain the solution [4]. Alternatively, a particle filter might be used, which avoids the linearisation of the measurement function [14]. However, in this case, naive application of a standard particle filter would be computationally infeasible because the number of basis functions used to describe the map, and hence the state dimension, could be very large depending on the length of pipe explored.

A further alternative is the Rao-Blackwellised particle filter (RBPF) [11], which in this case can be used to exploit the linear-in-the-parameters nature of the basis function decomposition of the map: a Kalman filter can be used to update the parameters $\boldsymbol{\theta}_k$, effectively reducing the state-dimension for the particle representation, keeping the computational complexity manageable. A particle filter can still be used to represent the robot location, which in this case is a one-dimensional problem, well suited to particle filtering. A similar formulation has been used in the popular FastSLAM algorithm [15].

The principle of Rao-Blackwellisation is to partition the joint distribution $p(\mathbf{x}_k, \boldsymbol{\theta} | \mathcal{Y}_k, \mathcal{U}_k, \mathbf{x}_0)$ using the product rule, where

$$p(\mathbf{x}_k, \boldsymbol{\theta}_k | \mathcal{Y}_k, \mathcal{U}_k, \mathbf{x}_0) = p(\boldsymbol{\theta}_k | \mathbf{x}_k, \mathcal{Y}_k) p(\mathbf{x}_k | \mathcal{Y}_k, \mathcal{U}_k, \mathbf{x}_0) \quad (10)$$

The key point to note is that the function $p(\boldsymbol{\theta}_k | \mathbf{x}_k, \mathcal{Y}_k)$ can be represented analytically, as a conditional Gaussian, so that only the function $p(\mathbf{x}_k | \mathcal{Y}_k, \mathcal{U}_k, \mathbf{x}_0)$ requires sampling [14]. This is why the map parameters $\boldsymbol{\theta}_k$ can be updated in a linear estimation step, i.e. by the Kalman filter.

A further significant point to note when using the RBPF algorithm in SLAM is that due to representing the robot pose via a set of particles, $\mathbf{x}_k^{(i)} \sim p(\mathbf{x}_k | \mathcal{Y}_k, \mathcal{U}_k, \mathbf{x}_0)$, for $i = 1, \dots, n_s$, where n_s is the number of particles, each individual particle has its own associated map. This means for every single particle representing robot pose $\mathbf{x}_k^{(i)}$, there is a distinct associated set of map parameters $\boldsymbol{\theta}_k^{(i)}$, and a separate Kalman filter must be computed to update each map.

C. PipeSLAM Algorithm

In this section the PipeSLAM algorithm is described for estimating the joint distribution $p(\mathbf{x}_k, \boldsymbol{\theta}_k | \mathcal{Y}_k, \mathcal{U}_k, \mathbf{x}_0)$ using the RBPF.

1) *Initialise particles:* Draw initial samples based on $\mathbf{x}_0^{(i)} \sim p(\mathbf{x}_0 | \mathcal{Y}_0)$, and initialise the parameter settings for the Kalman filter for all particle samples, such that $\boldsymbol{\theta}_0^{(i)} = \hat{\boldsymbol{\theta}}_0$, for $i = 1, \dots, n_s$, where n_s is the number of particle samples.

2) *Particle filter time update:* The particle state \mathbf{x}_k is updated at each time point according to linear state transition model (11),

$$\mathbf{x}_{k|k-1}^{(i)} = A_x \mathbf{x}_{k-1}^{(i)} + B_x \mathbf{u}_k + \boldsymbol{\omega}_k, \text{ for } i = 1, \dots, n_s \quad (11)$$

3) *Kalman filter time update:* The basis function weights for each map associated with the i^{th} particle, $\boldsymbol{\theta}_k^{(i)}$, for $i = 1, \dots, n_s$, are updated at time step k . Due to the assumption that the map is time-invariant, the weights $\boldsymbol{\theta}_k^{(i)}$ are also assumed to be unchanging over time, so the prediction step for $\boldsymbol{\theta}_k^{(i)}$ is a constant for the mean, and an inflation of the uncertainty in the covariance,

$$\boldsymbol{\theta}_{k|k-1}^{(i)} = \boldsymbol{\theta}_{k-1}^{(i)} \quad (12)$$

$$P_{k|k-1}^{(i)} = P_{k-1}^{(i)} + \Sigma_\eta \quad (13)$$

4) *Kalman filter measurement update:* The basis function weights for each map, $\boldsymbol{\theta}_k^{(i)}$, for $i = 1, \dots, n_s$, are corrected using the standard Kalman filter update step,

$$R_k^{(i)} = \Phi(\mathbf{x}_{k|k-1}^{(i)}) P_{k|k-1}^{(i)} \Phi^T(\mathbf{x}_{k|k-1}^{(i)}) + \Sigma_\nu \quad (14)$$

$$K_k^{(i)} = P_{k|k-1}^{(i)} \Phi^T(\mathbf{x}_{k|k-1}^{(i)}) R_k^{(i)-1} \quad (15)$$

$$\boldsymbol{\theta}_k^{(i)} = \boldsymbol{\theta}_{k|k-1}^{(i)} + K_k^{(i)} (\mathbf{y}_k - \Phi(\mathbf{x}_{k|k-1}^{(i)}) \boldsymbol{\theta}_{k|k-1}^{(i)}) \quad (16)$$

$$P_k^{(i)} = P_{k|k-1}^{(i)} - K_k^{(i)} R_k^{(i)} K_k^{(i)T} \quad (17)$$

5) *Evaluate and normalise importance weights:* Particle weights, $\alpha_k^{(i)}$ for $i = 1, \dots, n_s$, need to be re-evaluated based on the likelihood for the newly arrived observation, and then normalised,

$$\tilde{\alpha}_k^{(i)} = \alpha_{k-1}^{(i)} p(\mathbf{y}_k | \mathbf{x}_{k|k-1}^{(i)}, \boldsymbol{\theta}_{k|k-1}^{(i)}) \quad (18)$$

$$\alpha_k^{(i)} = \frac{\tilde{\alpha}_k^{(i)}}{\sum_{j=1}^{n_s} \tilde{\alpha}_k^{(j)}} \quad (19)$$

6) *Particle resampling:* When the number of effective particles (N_{eff}) drops to a certain threshold (N_{res}), the particle set needs to be resampled to avoid excessive particle depletion [14],

$$N_{\text{eff}} = \frac{1}{\sum_{j=1}^{n_s} (\alpha_k^{(j)})^2} \quad (20)$$

Therefore, when N_{eff} drops to a predefined threshold N_{res} , the particle set $\{\mathbf{x}_{k|k-1}^{(i)}, \boldsymbol{\theta}_k^{(i)}, P_k^{(i)}\}$ needs to be resampled by drawing n_s times from the present particle set to a new particle set and the probability of drawing a particle $\{\mathbf{x}_k^{(i)}, \boldsymbol{\theta}_k^{(i)}, P_k^{(i)}\}$ is proportional to its associated weight $\alpha_k^{(i)}$. The algorithm is then repeated from step 2.

D. Experimental details

In order to demonstrate the feasibility of the mapping and localisation technique described above, a small-scale laboratory experiment was constructed. A steel pipe, of dimensions 1 metre in length, by 88 millimetres in external diameter, was inserted into a concrete mould in a water butt, which was then filled with water (Fig. 2). The pipe material was chosen to be steel as opposed, to e.g. cast iron, because the acoustical properties of steel are well known. This makes steel a more appropriate material for testing in the early stages of developing this novel technique.

A 3D printed unit was used to house a pair of hydrophones (Bruel&Kjaer type 8103). This unit was then immersed into the steel pipe in an ultrasonic pulse-echo setup. The input signal to the pulser was a waveform generator (Tektronix AFG3022C), amplified by a Bruel&Kjaer type 2713 amplifier, which produced pipe vibration. The receiver unit, the second hydrophone connected to an additional amplifier (Bruel&Kjaer 2693), measured amplitude of the pipe vibration. Finally this output signal was logged on a PC using a National Instruments BNC 2110 receiving unit.

The map of pipe vibration amplitude (measured in arbitrary units, a. u.) was constructed over a 40 centimetre mid section of the pipe. Data was logged at 0.5 centimetre spacings over this mid section (see Fig. 2 for experimental data). A fast Fourier transform (FFT) was used to transform the data from the time- to frequency-domain to obtain the amplitude of pipe vibration at each spatial location. The amplitude was averaged over the range 15-25 kHz to produce a one-dimensional map of mean pipe vibration amplitude over space (Fig. 2(d)). This one-dimensional function in Fig. 2(d) corresponds to the map $h(\cdot)$ defined above in (2).

In this preliminary work the sensors were not mounted in the robot in Fig. 1, which would have overcomplicated the laboratory experiment. However, the robot is designed to carry a sensor payload such as used here.

E. Algorithm evaluation

The PipeSLAM algorithm was evaluated using a combination of experimental and simulation data. We used the map generated by the experiment described above as a ground truth map. We then simulated a robot moving back and forth along this map by the state-space model defined in (1) and (2).

The PipeSLAM algorithm was used in the simulation phase to simultaneously estimate both the map and robot location. Table I shows the parameter settings used in the PipeSLAM algorithm. In particular, the number of particle samples, n_s , was set to 100, which is relatively few but in this one-dimensional problem performed well.

To make the SLAM problem more challenging and highlight the benefit of using the map over dead reckoning, we also added a deterministic linear and sinusoidal drift term d_k to the state equation, of the form $d_k = -0.15\bar{m}_k + 0.02\bar{m}_k \sin(0.125\bar{m}_k)$, where $\bar{m}_k = km_k$, where $m_k = 0.0395$ cm.

TABLE I
PARAMETER SETTINGS FOR ALGORITHM EVALUATION

Parameter description	Parameters	Values
Number of basis functions	M	100
Width of basis functions	σ	1.5
State transition matrix of x	A_x	1
State transition matrix of θ	-	$\mathbf{I}_{100 \times 100}$
Input matrix of x	B_x	1
State noise covariance in x	Σ_ω	0.3 cm ²
State noise covariance in θ	Σ_η	$\mathbf{0}_{100 \times 100}$
Measurement noise covariance	Σ_ν	0.1
Particle sample size	n_s	100
Efficient particles threshold	N_{res}	50

TABLE II
ALGORITHM EVALUATION RESULTS

Estimation error description	Values
PipeSLAM location estimate RMSE	0.7426 cm
Dead reckoning location RMSE	3.0952 cm
PipeSLAM location estimate error $\sum_k e_k $	1279 cm
Dead reckoning location error $\sum_k e_k $	5713 cm
Final map RMSE	5.00 (a. u.)
Final map normalised RMSE	0.04 (a. u.)
Final map error $\sum_k e_k $	5093 (a. u.)

III. RESULTS AND DISCUSSION

In order to test and evaluate the PipeSLAM algorithm, we used the experimental pipe vibration data to define the ground truth map (Fig.2). We then simulated the robot moving forward and backward over this data to analyse the PipeSLAM algorithm. Drift was applied to the robot simulation to make the SLAM problem more challenging and the map used in the PipeSLAM algorithm was initialised with zero weights.

The results obtained from the simulation demonstrate how the PipeSLAM algorithm learns the map: each particle used to represent robot location has an associated map based on the trajectories of the particles - these separate maps are clearly seen in Fig. 3(a)-(b). On the backward pass along the simulated pipe the maps have converged to similar values (Fig. 3(c)-(d)). This convergence of the maps is due to resampling the particles, which leads to reduced particle diversity, and hence reduced map diversity, which is a common issue in particle filter based SLAM [15].

Regarding localisation, on the forward pass, whilst the map is unknown there is relatively large uncertainty in the robot location. Then on the backward pass the uncertainty reduces (Fig.3(e)). Quantitatively, the PipeSLAM algorithm was able to localise more accurately than dead reckoning, in particular, the cumulative error using PipeSLAM was $\sum_k |e_k| = 1279$ cm and using dead reckoning was $\sum_k |e_k| = 5713$ cm (see Table II), which demonstrates the effectiveness of the PipeSLAM algorithm.




Fig. 2. Experimental setup. (a) A one metre steel pipe was immersed in water and a hydrophone pulser and receiver unit, H8103, was moved up and down the pipe to excite and record pipe vibration. The experimental data used here was from a 40 cm mid section of pipe. (b) Signals measured by the hydrophone over the 40 cm mid section of pipe. Signals were observed every 0.5 cm but for clarity the graph only shows signals at 5 cm spacings. (c) The space-frequency representation of pipe vibration amplitude. The amplitudes over frequency were obtained from an FFT of the time-domain signals. The red lines define the frequency range used to average over to form the map. (d) The map of pipe vibration amplitude over space.

IV. CONCLUSIONS

In this paper we have addressed the SLAM problem for feature-sparse water pipes. To overcome the problem of feature-sparsity, we have developed a technique of exciting and measuring pipe vibration using a pair of hydrophones to create a map of mean vibration amplitude along the length of the pipe. In order to use this map we have developed a PipeSLAM algorithm based on the Rao-Blackwellised particle filter: the key features of this algorithm are that the map is decomposed using a radial basis function network that is linear-in-the-parameters; these map parameters are estimated using a Kalman filter; robot location is estimated using particles. We have tested the feasibility of the algorithm using a combination of experiment and simulation, where the map of pipe vibration amplitude was generated in experiment and then the PipeSLAM algorithm was evaluated on the map in simulation. A benchmark against robot localisation using just dead reckoning demonstrated the success of the PipeSLAM method.

REFERENCES

- [1] Z. Liu and Y. Kleiner, "State of the art review of inspection technologies for condition assessment of water pipes," *Measurement*, vol. 46, no. 1, pp. 1–15, 2013.
- [2] J. M. Mirats Tur and W. Garthwaite, "Robotic devices for water main in-pipe inspection: A survey," *Journal of Field Robotics*, vol. 27, no. 4, pp. 491–508, 2010.
- [3] J. Moraleda, A. Ollero, and M. Orte, "A robotic system for internal inspection of water pipelines," *IEEE Robotics & Automation Magazine*, vol. 6, no. 3, pp. 30–41, 1999.
- [4] H. Durrant-Whyte and T. Bailey, "Simultaneous localization and mapping: part I," *IEEE Robotics & Automation Magazine*, vol. 13, no. 2, pp. 99–110, 2006.
- [5] D. Kryz and H. Najjaran, "Development of visual simultaneous localization and mapping (VSLAM) for a pipe inspection robot," in *International Symposium on Computational Intelligence in Robotics and Automation*, 2007, pp. 344–349.
- [6] H. Lim, J. Y. Choi, Y. S. Kwon, E.-J. Jung, and B.-J. Yi, "SLAM in indoor pipelines with 15mm diameter," in *IEEE International Conference on Robotics and Automation*, 2008, pp. 4005–4011.
- [7] K. Ma, J. Zhu, T. J. Dodd, R. Collins, and S. R. Anderson, "Robot mapping and localisation for feature sparse water pipes using voids as landmarks," in *Towards Autonomous Robotic Systems: 16th Annual Conference, TAROS 2015*, Liverpool, UK, 2015, pp. 161–166.
- [8] L. D. Hostetler and R. D. Andreas, "Nonlinear Kalman filtering techniques for terrain-aided navigation," *IEEE Transactions on Automatic Control*, vol. 28, no. 3, pp. 315–323, 1983.
- [9] K. Ma, M. Schirru, A. H. Zahraee, R. Dwyer-Joyce, J. Boxall, T. J. Dodd, R. Collins, and S. R. Anderson, "Robot mapping and localisation in metal water pipes using hydrophone induced vibration and map alignment by dynamic time warping," in *IEEE International Conference on Robotics and Automation*, 2017, accepted.
- [10] S. Barkby, S. B. Williams, O. Pizarro, and M. V. Jakuba, "Bathymetric particle filter SLAM using trajectory maps," *The International Journal of Robotics Research*, vol. 31, no. 12, pp. 1409–1430, 2012.
- [11] A. Doucet, N. de Freitas, and N. Gordon, *An Introduction to Sequential Monte Carlo Methods*. New York, NY: Springer New York, 2001.
- [12] D. D. Ratnayaka, M. J. Brandt, and K. M. Johnson, *Twort's Water Supply*, 6th ed. Oxford, UK: Butterworth-Heinemann, 2009.
- [13] V. Kadirkamanathan and M. Niranjan, "A function estimation approach to sequential learning with neural networks," *Neural Computation*, vol. 5, no. 6, pp. 954–975, 1993.
- [14] S. Sarkka, *Bayesian Filtering and Smoothing*. Cambridge, UK: Cambridge University Press, 2013.
- [15] M. Montemerlo and S. Thrun, *FastSLAM: A Scalable Method for the Simultaneous Localization and Mapping Problem in Robotics*. Berlin, Heidelberg: Springer, 2007.

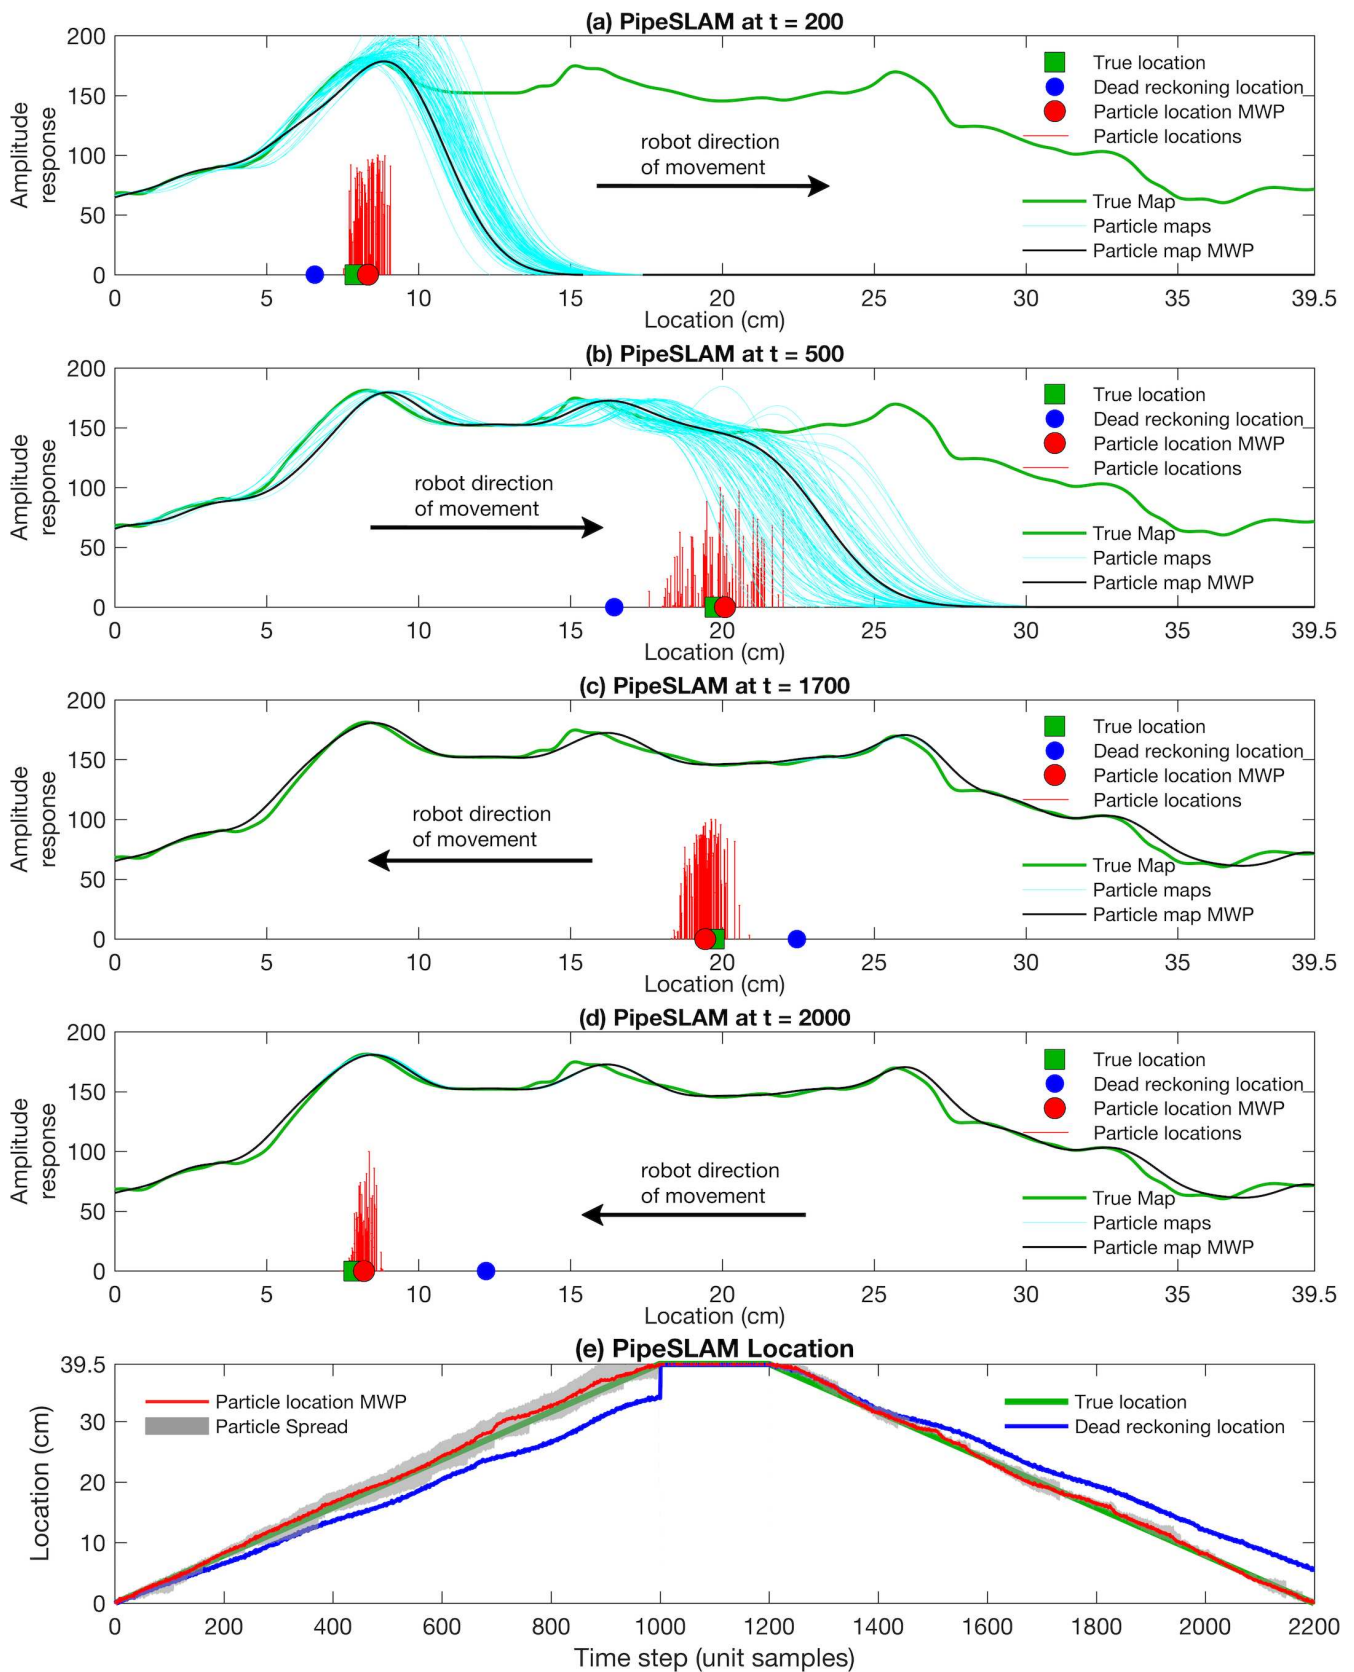


Fig. 3. PipeSLAM results from simulation of robot movement along a pipe, forwards and backwards, over a ground truth map generated from experimental hydrophone-induced pipe vibration data. (a)-(d) PipeSLAM at increasing time steps of robot movement (panels a-b: forward movement, left to right; panels c-d: backward movement, right to left). Note that MWP refers to the maximum weighted particle. (e) PipeSLAM localisation results over time, showing the improvement compared to the dead reckoning estimate, which exhibits significant drift.