



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/112864/>

Version: Accepted Version

Proceedings Paper:

Ma, K., Schirru, M.M., Zahraee, A.H. et al. (2017) 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 IEEE International Conference on Robotics and Automation, 29/05/2017 - 03/06/2017, Singapore. Institute of Electrical and Electronics Engineers, pp. 2548-2553. ISBN: 978-1-5090-4633-1. ISSN: 1050-4729.

<https://doi.org/10.1109/ICRA.2017.7989296>

© 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

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

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.

Robot mapping and localisation in metal water pipes using hydrophone induced vibration and map alignment by dynamic time warping*

Ke Ma¹, Michele M. Schirru², Ali Hassan Zahraee¹, Rob Dwyer-Joyce², Joby Boxall³, Tony J. Dodd¹, Richard Collins³ and Sean R. Anderson¹

Abstract—Water is a highly valuable resource so asset management of associated infrastructure is of critical importance. Water distribution pipe networks are usually buried, and so are difficult to access. Robots are therefore appealing for performing inspection and detecting damage to target repairs. However, robot mapping and localisation of buried water pipes has not been widely investigated to date, and is challenging because pipes tend to be relatively featureless. In this paper we propose a mapping and localisation algorithm for metal water pipes with two key novelties: the development of a new type of map based on hydrophone induced vibration signals of metal pipes, and a mapping algorithm based on spatial warping and averaging of dead reckoning signals used to calibrate the map (using dynamic time warping). Localisation is performed using both terrain-based extended Kalman filtering and also particle filtering. We successfully demonstrate and evaluate the approach on a combination of experimental and simulation data, showing improved localisation compared to dead reckoning.

I. INTRODUCTION

Water is one of our most valuable natural resources, and so management of associated infrastructure assets is of critical importance. Water is usually distributed by underground pipe networks, which are difficult for humans to access. So there is a significant amount of interest in developing technologies that can inspect, maintain and repair buried water pipes [1], [2]. Robots would appear to have great potential for inspecting these difficult-to-access pipe networks [3], [4]. However, whilst there has been much research devoted to designing robots and sensors for damage detection in water pipes [3], mapping and localisation has received less attention. Accurate mapping and localisation is of key importance for precisely locating areas of poor condition and damage, and therefore targeting maintenance action from above ground, to minimise costs and disruption to the pipe network and water supply. This mapping and localisation problem is addressed here in the context of metal water pipes, focusing on ways of transforming the apparently featureless environment of the water pipe, into a more feature-rich environment.

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

¹Ke Ma, Tony J. Dodd, Ali Hassan Zahraee and Sean R. Anderson are with the Department of Automatic Control and Systems Engineering, University of Sheffield, Sheffield, S1 3JD, UK ke.ma t.j.dodd a.hassanzahraee s.anderson@sheffield.ac.uk

²Michele Schirru, and Rob Dwyer-Joyce are with the Department of Mechanical Engineering, University of Sheffield, Sheffield, S1 3JD, UK m.schirru r.dwyer-joyce@sheffield.ac.uk

³Joby Boxall and Richard Collins are with the Department of Civil and Structural Engineering, University of Sheffield, Sheffield, S1 3JD, UK j.b.boxall r.p.collins@sheffield.ac.uk

In previous work we have highlighted a number of particular challenges for robot mapping and localisation in water pipes [5]. The first is that the water pipe is a relatively feature sparse environment compared to typical robot mapping and localisation scenarios. The second challenge is that typical robotics sensors used for range and bearing can only detect nearby features, especially in pipes of small diameter (~ 3 inches), which are commonly used for distribution to buildings off the larger trunk mains. The third challenge is that the robot has restricted movement, either forwards or backwards along the pipe, and cannot survey the environment from different perspectives to reduce uncertainty in both the map and the robot location. Therefore, robot mapping and localisation in water pipes presents a variety of unique and challenging problems.

Simultaneous localisation and mapping (SLAM) techniques [6] have been applied to robots in water pipes, using cameras and inertial measurement units (IMUs) [7], [8]. However, the use of cameras is limited by the lack of visual features, and IMUs are subject to drift, meaning that there is still much scope for improvement. In previous work, we have proposed and demonstrated the idea of using ultrasonic transceivers to sense soil depth through plastic pipe walls, and used this map in a type of terrain-based localisation scheme [5]. In that work we did not solve the problem of constructing a calibrated spatial map from dead reckoning sensors under drift conditions, nor develop a similar technique suited to metal pipes. We address both of those issues here.

One of the key novel ideas in this work is to create a new type of map based on signals from hydrophone excitation of the metal pipe. This leads to a map of pipe vibration amplitude over space. The problem is that the only way of calibrating the spatial component of the map is by a dead reckoning sensor such as a motor encoder. We assume that the robot will make multiple passes up and down the pipe between two known locations. This means that the map calibration can be improved by spatial averaging. However, taking a direct average of the data would be likely to lead to smoothing of the peaks and troughs in the map, degrading features required for localisation. Instead, the second key novel idea we propose is to use a signal alignment technique to warp the maps in the spatial direction before averaging. This improves the spatial calibration without overly smoothing the map features. The signal alignment and averaging algorithm is based on dynamic time warping (DTW) and is known as DTW barycentre averaging (DBA) [9].

The paper is structured as follows. In section II we define the mapping and localisation problem, present the mapping solutions based on DBA and the two alternatives to localisation, one based on the extended Kalman filter (EKF) used in terrain-based localisation [10], and one based on the particle filter (PF) [11]. The experimental details for the hydrophone induced vibration of the metal pipe is also presented in this section. In section III we present the results on evaluation of the mapping and localisation methods using a combination of experimental and simulation data. Finally, in section IV we conclude the paper.

II. METHODS

A. Problem statement

In defining the concept of operations for the pipe inspection robot considered in this work (see Fig. 1 for a prototype), we make a number of assumptions. The first regards the robot deployment. In consultation with project partners from the water utilities industry, e.g. Yorkshire Water, we require the robot to enter the water pipe network through existing access points to minimise costs. We assume that fire hydrants could be used, which in the UK, Europe and the USA are separated by approximately 100 metres [12, Chapter 14, Table 14.2].

This leads to the second assumption, that the robot will travel between two points with known location, i.e. two fire hydrants. The third assumption is that robot travel time between two hydrants, in terms of relative time cost to e.g. robot deployment, is relatively trivial, hence it is worth the robot making multiple passes up and down the pipe in order to maximise the mapping and localisation accuracy.

Due to the fact that navigation through the pipe in itself is relatively trivial, i.e. forwards or backwards, the need for a SLAM solution is limited, hence, the sequential approach to mapping and localisation taken here. In addition, we only consider the problem of localising with respect to one dimension, i.e. distance travelled through the pipe: this technique is not suited to correcting heading, which we leave to future work.

We define the mapping and localisation problem as follows: 1. to estimate a map $h(\mathbf{x}_k)$ from hydrophone-induced pipe vibration signals, that transforms robot pose $\mathbf{x}_k \in \mathbb{R}^{n_x}$ at time-step k to sensor measurements $\mathbf{y}_k \in \mathbb{R}^{n_y}$, where $h: \mathbf{x}_k \rightarrow \tilde{\mathbf{y}}_k$, and robot pose for example is $\mathbf{x}_k = [x \ y \ \theta]^T$, i.e. location in x - y co-ordinates and heading θ , where $\tilde{\mathbf{y}}_k$ is the noise-free sensor output; and 2. localise the robot by obtaining the estimate of the pose distribution $p(\mathbf{x}_k|\mathbf{y}_k)$.

We assume that the dynamics of the pipe robot can be represented by a state-space model, with state dynamics

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

where $g(\cdot)$ is the state transition function, $\mathbf{u}_k \in \mathbb{R}^{n_u}$ is the input, $\mathbf{w}_k \sim N(\mathbf{0}, \Sigma_w)$ is the state noise. The measurement model is

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

$h(\cdot)$ is the measurement function, and $\mathbf{v}_k \sim N(\mathbf{0}, \Sigma_v)$ is the measurement noise.

In this paper we simplify the state vector \mathbf{x}_k to contain just the location of the robot along the pipe, $\mathbf{x}_k = x_k$ and $n_x = 1$. The observation \mathbf{y}_k is the processed hydrophone signal, which is the average of the vibration amplitude over some frequency range, $|\bar{a}|_k$, hence, $\mathbf{y}_k = |\bar{a}|_k$, and $n_y = 1$. The state dynamics are assumed to be obtainable from a processed motor encoder reading, which defines distance travelled, m_k , hence $g(\mathbf{x}_{k-1}, \mathbf{u}_{k-1}) = A\mathbf{x}_{k-1} + B\mathbf{u}_{k-1}$, where $A = 1$, $B = 1$, and $\mathbf{u}_{k-1} = m_k$. Although defining this one dimensional state-space model is relatively trivial, it has the advantage that it provides ready extensibility to more state dimensions for representing the pose in two or three dimensions, and also can incorporate more sensors, e.g. camera and IMU data.

B. Map construction using signal alignment algorithm DBA

We demonstrate in this work that the robot can obtain a map of pipe vibration amplitude over space by travelling through the pipe and exciting pipe vibration using a hydrophone. Corresponding locations of the robot can be calibrated using dead reckoning, e.g. from a motor encoder. However, any drift in the dead reckoning estimate will result in an incorrectly spatially calibrated map. A solution to this problem is for the robot to make multiple passes back and forth through the pipe in order to generate a set S of n_m independent sequences of map data, in order to average out drift errors, where

$$S = \{\mathbf{s}_1, \dots, \mathbf{s}_{n_m}\} \quad (3)$$

where each data sequence is comprised of pipe amplitude response signals over space, $\mathbf{s}_j = (|\bar{a}|_{j,1}, \dots, |\bar{a}|_{j,n_s})$, where n_s is the number of spatial samples, and for each amplitude datum there is a corresponding dead reckoning estimate of spatial location, $x^{(d)}$, hence we have the data pairs $(x_{j,k}^{(d)}, |\bar{a}|_{j,k})$, for each map $j = 1, \dots, n_m$ and for each observation within the map $k = 1, \dots, n_s$.

The sequences in the set S can be combined to reduce the effect of drift, however, a direct averaging of these sequences would be likely to smooth out the map due to sequence misalignment, degrading features required for localisation. Instead, we propose that sequences can be combined into a map using a signal alignment technique that warps the sequences in the spatial dimension before averaging. In this work we use a signal alignment technique based on dynamic time warping (DTW), known as DTW barycentre averaging (DBA).

The DTW algorithm calculates an alignment cost matrix M , between two sequences \mathbf{s}_k and \mathbf{s}_l (see Algorithm 1 in Fig. 3). The optimal alignment between the sequences follows a ‘valley’ in the cost matrix. The approach of DBA is to use DTW to compare a mean sequence estimate, $\bar{\mathbf{s}}$, to each sequence in S and iteratively reduce the total DTW cost,

$$\bar{\mathbf{s}}^* = \arg \min_{\bar{\mathbf{s}}} \sum_{i=1}^{n_m} D^2(\bar{\mathbf{s}}, \mathbf{s}_i) \quad (4)$$

where the quantity $D(\bar{\mathbf{s}}, \mathbf{s}_i)$ is the cumulative alignment cost calculated by DTW, where $D(\bar{\mathbf{s}}, \mathbf{s}_i) = M_{N_1, N_2}$ is obtained




Fig. 1. Design of a water pipe inspection robot for small diameter water pipes. The prototype shown here is designed to operate in 3 inch pipe and is tethered to simplify recovery in the event of a robot failure (a requirement of water utilities). The use of a tether can also be exploited to supply power and off-board processing. The robot is composed of modules that are flexibly linked by a steel spring. Each module is composed of a 70mm long by 29 mm diameter core unit, which contains sensors and processing units. The flexible arms extend the diameter of each module to the range 65-80mm, bracing against the inner pipe wall for stability. (a) Robot with 3 modules. (b) Robot in 3 inch clear plastic pipe. (c) Zoomed view of one robot module.

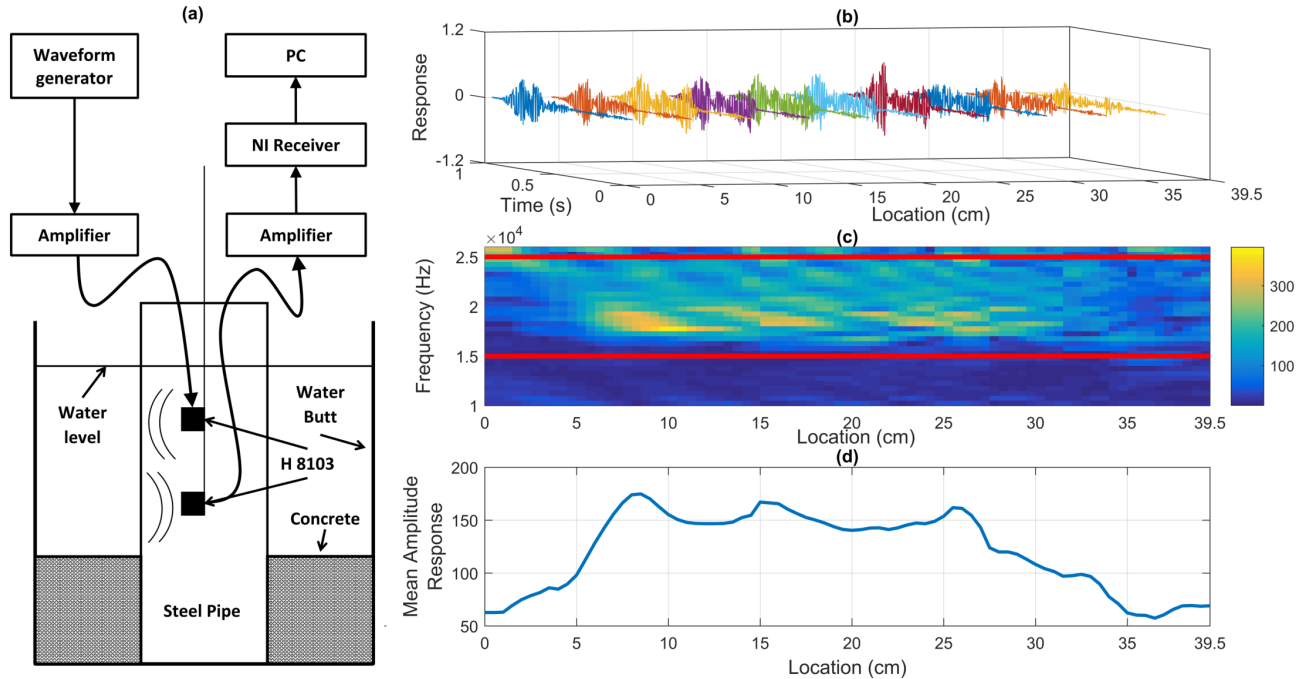


Fig. 2. Experimental setup. (a): The hydrophone pulser and receiver unit, H8103, travels up and down a one metre steel pipe that is immersed in water. Experimental recording was conducted on a 40 cm mid-section of pipe. (b): Time-domain signals observed by the hydrophone over 40 cm of pipe. Signal were observed at 0.5 cm intervals but for clarity the graph only shows signals at 5 cm spacing. (c): Space-frequency representation of hydrophone signal amplitude obtained from an FFT of the time-domain signals. The red lines indicate the region over which the average amplitude is taken to form the one dimensional map. (d): The hydrophone map of amplitude over space.

from the final element computed for the cost matrix M in Algorithm 1 (Fig. 3).

The algorithm DBA can be used to obtain a solution to the optimisation problem posed in (4), i.e. the optimal signal average, \bar{s}^* . The algorithm is iterative and has guaranteed convergence [13]. At each iteration k :

- 1) Use DTW (Algorithm 1) to iteratively compute the optimal alignment between each data sequence and the current estimate of the signal average \bar{s}_k , i.e. $\text{DTW}(\bar{s}_k, s_j)$, for $j = 1, \dots, n_m$.
- 2) Use the updated alignment from step 1 to update the signal average to \bar{s}_{k+1} and set $\bar{s}^* = \bar{s}_{k+1}$. Increment k and go to step 1.

The initial mean \bar{s}_0 is defined by using one of the data sequences in S chosen at random. The algorithm is repeated until convergence, which can be monitored by evaluating the cumulative alignment cost in (4).

Finally, the optimal sequence of data samples \bar{s}^* forms the continuous map function $h(\cdot)$ from linear interpolation of the data pairs $(x_1^*, s_1^*), (x_2^*, s_2^*), \dots, (x_{n_s}^*, s_{n_s}^*)$, where for a location x_k on the interval (x_j, x_{j+1}) , at sample-time k , we define

$$h(x_k) = s_j^* + (s_{j+1}^* - s_j^*) \frac{x_k - x_j^*}{x_{j+1}^* - x_j^*}. \quad (5)$$

Here we use linear interpolation to define the map $h(x_k)$ but an alternative such as splines could equally be used.

Algorithm 1: DTW

```

1: procedure DTW( $\mathbf{s}_k, \mathbf{s}_l$ )
2:    $M_{1,1} = 0$ 
3:    $M_{2:N_k,1} = \infty$ 
4:    $M_{1,2:N_l} = \infty$ 
5:   for  $i = 2$  to  $N_k$  do
6:     for  $j = 2$  to  $N_l$  do
7:        $c = d(\mathbf{s}_k(i), \mathbf{s}_l(j))$ 
8:        $M_{i,j} = c + \min(M_{i-1,j}, M_{i,j-1}, M_{i-1,j-1})$ 
9:     end for
10:  end for
11: end procedure

```

Fig. 3. Algorithm 1: The dynamic time warping (DTW) algorithm, which takes as input two sequences of data, $\mathbf{s}_k \in \mathbb{R}^{N_k}$ and $\mathbf{s}_l \in \mathbb{R}^{N_l}$, and returns the matrix of similarity measure between the two sequences $M \in \mathbb{R}^{N_k \times N_l}$.

C. Localisation by extended Kalman filtering

The approach taken to localisation using the EKF is inspired by a terrain-based navigation algorithm developed for aerospace applications [10]. The steps for the EKF at sample time k consist of:

- 1) The prediction step for the state vector x_k and state covariance P_k ,

$$x_k^- = Ax_{k-1} + Bu_{k-1} \quad (6)$$

$$P_k^- = AP_{k-1}A^T + Q \quad (7)$$

where $Q = \Sigma_w$ is the state noise covariance.

- 2) The measurement update step requires the definition of the linearised measurement model, H_k , which is obtained from the derivative of a local quadratic fit to the spatial map of pipe vibration amplitude, where the local quadratic approximation of the spatial map is

$$h(x_k) = ax_k^2 + bx_k + c \quad (8)$$

where x_k is the current location of the robot and hence the derivative is

$$H_k = \frac{d}{dx_k} h(x_k) = 2ax_k + b \quad (9)$$

The parameters of the local quadratic function are obtained from a least-squares fit to a data window centred on the current estimate of the robot location x_k : the size of local quadratic fit window was set proportional to the state covariance, $W_{fit} = \alpha P_k$ [10].

- 3) The EKF measurement update is performed by

$$K_k = P_k^- H_k^T (H_k P_k^- H_k^T + R + \epsilon_k^2)^{-1} \quad (10)$$

$$\hat{x}_k = x_k^- + K_k (y_k - h(x_k^-)) \quad (11)$$

$$P_k = (I - K_k H_k) P_k^- \quad (12)$$

where $R = \Sigma_v$ is the measurement noise covariance, and the term ϵ_k^2 , in (10), is due to the linear fit error ϵ_k , where

$$\epsilon_k = h(x_k) - 2ax_k - b \quad (13)$$

This time-varying error term is recommended by [10] because it inflates the measurement noise covariance

term R_k in regions of poor fit (typically due to high nonlinearity), reducing the chance of filter divergence and making the EKF more robust.

D. Localisation by particle filtering

As an alternative to the EKF for localisation we also investigated the use of the PF. In this case the PF is based on the bootstrap filter with sequential importance resampling [11] as used in our previous work on mapping and localisation in plastic water pipes [5].

In the first step the particles are initialised,

$$x_0^{(i)} \sim p(x_0), \text{ for } i = 1, \dots, n_s \quad (14)$$

where n_s is the number of samples, and initial particle weights are set to $w_0^{(i)} = \frac{1}{n_s}$, for $i = 1, \dots, n_s$. At each sample-time k the PF performs the following steps:

- 1) The location is predicted by samples drawn from the state equation, Eq. 1,

$$x_k^{(i)} \sim p(x_k | x_{k-1}^{(i)}, u_{k-1}), \text{ for } i = 1, \dots, n_s \quad (15)$$

where we assume that the state equation can be used as the importance distribution of the particle filter [11].

- 2) The weight update step is

$$w_k^{(i)} \propto p(y_k | x_k^{(i)}), \text{ for } i = 1, \dots, n_s \quad (16)$$

where we assume Gaussian noise \mathbf{v}_k on the sensor output,

$$w_k^{(i)} = \exp\left(-\frac{1}{2} \left(y_k - \hat{y}_k^{(i)}\right)^T R^{-1} \left(y_k - \hat{y}_k^{(i)}\right)\right), \quad (17)$$

for $i = 1, \dots, n_s$, where $\hat{y}_k^{(i)} = h(x_k^{(i)})$. The weights are then normalised to sum to unity.

- 3) To avoid degeneracy, resampling is performed if the effective number of particles drops below a threshold, $\gamma = 0.6n_s$, using stratified resampling [11].

E. Experimental details

A steel pipe (1m length, 88 mm in external diameter) was constrained by a concrete mould in a water butt that was filled with water (Fig. 2). Steel was chosen instead of cast iron because the acoustical properties of steel are well known, making it the ideal material for testing this new technique.

A pair of hydrophones (Bruel&Kjaer type 8103) were immersed into the pipe in an ultrasonic pulse-echo setup, contained in a 3D printed unit, of length 100 mm. In this setup, an arbitrary waveform generator (Tektronix AFG3022C) was used to excite with an electric signal one hydrophone, the pulser, that had the function to produce the ultrasonic vibration. This signal was amplified by Bruel&Kjaer type 2713 amplifier. The second hydrophone, the receiver, had the function to record the amplitude of the pipe vibration, similarly to a microphone when recording a sound pattern. The receiver was connected to a second amplifier (Bruel&Kjaer 2693) and the received signal was logged in a computer using a NI BNC 2110 receiving unit.

The hydrophones were submerged through depths of 0.5 cm spacing to observe the vibration signal obtained over a 40 cm mid section of the pipe (see Fig. 2 for experimental data). The time-domain data at each spatial location was transformed to the frequency-domain using a fast Fourier transform (FFT). Then the amplitude at each location was averaged across the frequency range 15-25 kHz to obtain the map data (Fig. 2).

In this investigation the hydrophones were not mounted on the robot shown in Fig. 1, which would have been an unnecessary complication for this work, but the pulse-receiver unit is suited to mounting on that robot design.

F. Algorithm evaluation

The mapping and localisation algorithms were evaluated by combining the data from the hydrophone experiment with a simulated robot moving up and down the map, using the state-space model defined in (1) and (2) for the simulation. The number of map sequences generated was $n_m = 20$, five passes forward along the pipe and five backwards. The input was constant, $m_k = 0.1$ cm or $m_k = -0.1$ cm depending on direction, and drift was added to the simulated robot in the mapping stage in the form of white noise \mathbf{w}_k , i.e. state noise covariance $\Sigma_w = 0.05$ cm², also the measurement noise term \mathbf{v}_k was set to $\Sigma_v = 0.01$.

For the EKF localisation the size of quadratic fit window was set to $W_{fit} = \alpha P_k$, with $\alpha = 2.3$; the noise covariances were set to $\sqrt{Q} = 0.5$ cm and $\sqrt{R} = 5$. For the particle filter localisation, the number of particles was set to $n_s = 300$ and noise covariances were set to $\sqrt{Q} = 0.5$ cm and $\sqrt{R} = 5$. To make the localisation 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.

III. RESULTS AND DISCUSSION

In order to evaluate the mapping and localisation algorithms described above we used the experimental data to define a ground truth map. We then simulated robot movement up and down this map, with simulated drift, to investigate the effectiveness of the DBA algorithm for constructing the map: we generated a total of 20 maps to align and average using DBA (Fig. 4(a)). We then applied both the EKF and the PF localisation algorithms to this estimated map and found that the technique improved on using dead reckoning alone as expected (Fig. 4(b)-(e)): the sum of absolute errors for dead reckoning was $\sum_k |e_k| = 5706$, for the EKF was $\sum_k |e_k| = 975$, and for the PF was $\sum_k |e_k| = 1046$. Hence, both the EKF and PF greatly outperformed the localisation using dead reckoning alone. The EKF slightly outperformed the PF, which along with the efficiency of the EKF approach, makes the EKF more appealing in this application.

These results give the first evidence that mapping and localisation using hydrophone induced vibration with map alignment using DBA is feasible, supporting field testing of the sensor on the robot prototype shown in Fig. 1(a). The

technique for calibrating the spatial map is also extensible to other types of sensor that would produce similar map data, e.g. through-pipe-wall ultrasonics, proposed in our earlier work [5]. One appealing feature for the hydrophone method we test here, in comparison to the through-pipe-wall ultrasonic method, is that it is omnidirectional in nature, which should make it robust to any robot rotations. A limitation of the method is that it is only useful for correcting drift along the length of the pipe, i.e. distance travelled, not heading estimates. This we leave to future work, but envisage fusing the method with IMU data to solve this problem.

IV. CONCLUSIONS

In this paper we have addressed the problem of robot mapping and localisation in metal water pipes. We have developed a novel technique for constructing a map using hydrophone induced pipe vibration. In order to address the problem of spatially calibrating the map using a dead reckoning sensor subject to drift, we proposed the use of a signal alignment and averaging algorithm based on dynamic time warping. The localisation was based on nonlinear state estimation. We evaluated the approach on a combination of experimental and simulation data, demonstrating that the technique is effective.

REFERENCES

- [1] T. Hao, C. Rogers, N. Metje, D. Chapman, J. Muggleton, K. Foo, P. Wang, S. R. Pennock, P. Atkins, S. Swingle *et al.*, "Condition assessment of the buried utility service infrastructure," *Tunnelling and Underground Space Technology*, vol. 28, pp. 331-344, 2012.
- [2] 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.
- [3] 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.
- [4] 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.
- [5] 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.
- [6] H. Durrant-Whyte and T. Bailey, "Simultaneous localization and mapping: part I," *IEEE Robotics & Automation Magazine*, vol. 13, no. 2, pp. 99-110, 2006.
- [7] D. Krysz 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. CIRA 2007. IEEE, 2007*, pp. 344-349.
- [8] 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. ICRA 2008., 2008*, pp. 4005-4011.
- [9] F. Petitjean, A. Ketterlin, and P. Gançarski, "A global averaging method for dynamic time warping, with applications to clustering," *Pattern Recognition*, vol. 44, no. 3, pp. 678-693, 2011.
- [10] 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.
- [11] S. Sarkka, *Bayesian Filtering and Smoothing*. Cambridge, UK: Cambridge University Press, 2013.
- [12] D. D. Ratnayaka, M. J. Brandt, and K. M. Johnson, *Twort's Water Supply*, 6th ed. Oxford, UK: Butterworth-Heinemann, 2009.
- [13] F. Petitjean, G. Forestier, G. I. Webb, A. E. Nicholson, Y. Chen, and E. Keogh, "Dynamic time warping averaging of time series allows faster and more accurate classification," in *2014 IEEE International Conference on Data Mining. IEEE, 2014*, pp. 470-479.

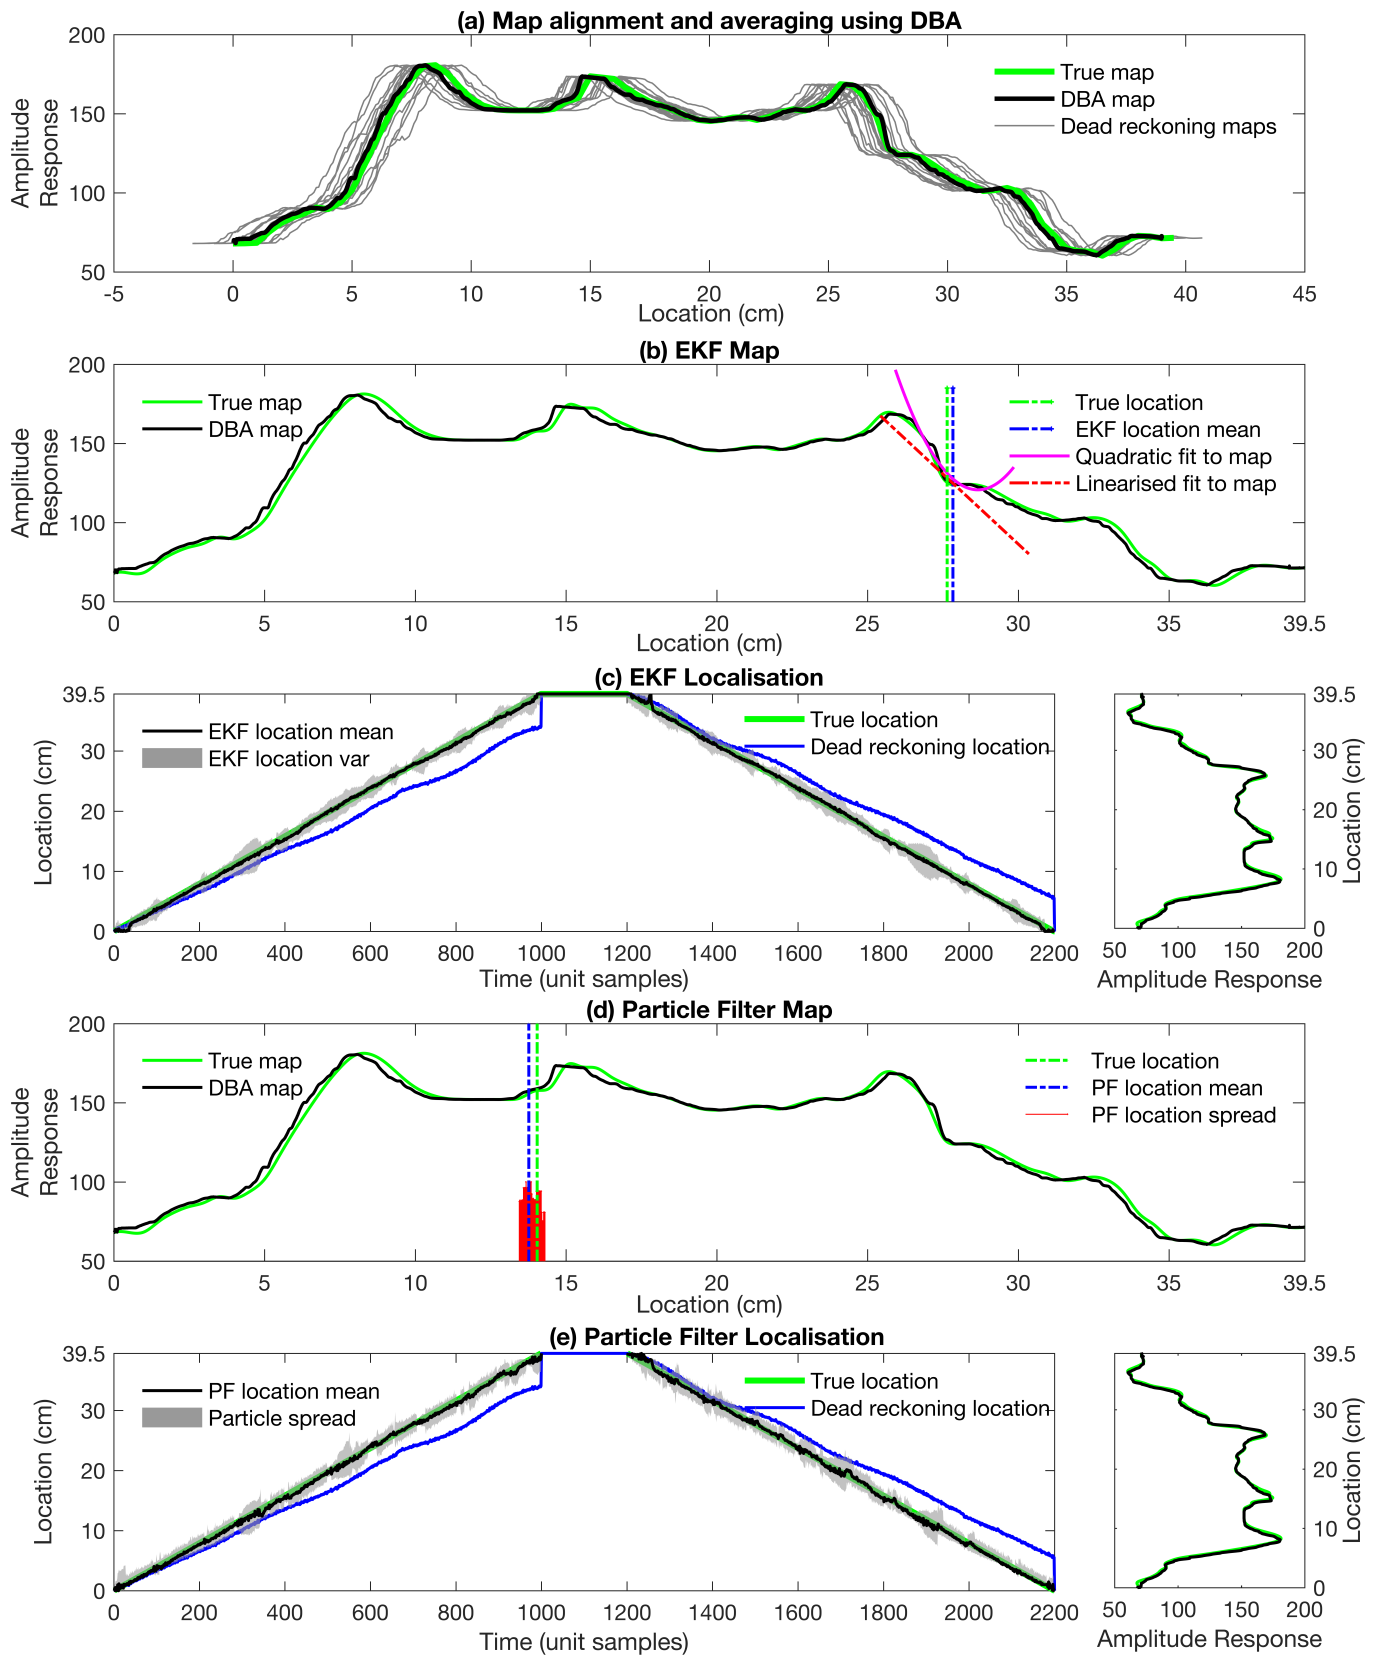


Fig. 4. Mapping and localisation results. (a) The use of DBA to construct a map estimate from observations with simulated drift. (b) and (c) Localisation using an extended Kalman filter (EKF). (d) and (e) Localisation using a particle filter (PF). For both the EKF and PF a comparison is given to dead reckoning, showing the clear improvement in localisation accuracy with EKF and PF.