



Towards ultrasound-based localization and mapping for long-range inspection robots

Othmane-Latif Ouabi

► To cite this version:

Othmane-Latif Ouabi. Towards ultrasound-based localization and mapping for long-range inspection robots. Robotics [cs.RO]. Université de Lorraine; Georgia Tech-Lorraine, 2022. English. NNT: . tel-03835195

HAL Id: tel-03835195

<https://hal.science/tel-03835195>

Submitted on 31 Oct 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Towards ultrasound-based localization and mapping for long-range inspection robots

THÈSE

présentée et soutenue publiquement le 10 Octobre 2022

pour l'obtention du

Doctorat de l'Université de Lorraine

(mention Automatique et Traitement du signal)

par

Othmane-Latif Ouabi

Composition du jury

<i>Président :</i>	Nico F. Declercq
<i>Directeur de thèse :</i>	Cédric Pradalier
<i>Co-directeur de thèse :</i>	Matthieu Geist
<i>Co-encadrant :</i>	Nico F. Declercq
<i>Rapporteurs :</i>	Teresa Vidal Calleja Philippe Lasaygues
<i>Examineurs :</i>	Carol Martinez Luna
<i>Membres invités :</i>	Neil Zeghidour Henri Walaszek

Mis en page avec la classe thesul.

Acknowledgments

First and foremost, I would like to thank my supervisors, Cédric Pradalier and Matthieu Geist, who gave me the opportunity to do a PhD. It is incredible how much I've learned from them. I am grateful for their sustained and unconditional support, over the past few years. My sincere thanks shall also go to Nico F. Declercq, Neil Zeghidour, and Pascal Pomarede, with whom it was always pleasant to discuss, whether it be on scientific or non-scientific matters. They have undoubtedly been a huge support.

I really appreciated the time I spent with all my colleagues, both in and outside the lab. I will remember the badminton nights with Assia Benbihi, Antoine Mahé, and Laura Monnier, that were abruptly put to an end due to a well-known pandemic. I won't forget the nice moments I shared with Antoine Richard and Stéphanie Aravecchia, along with their guidance when I needed to do either Linux-related or manual things. I appreciated my discussions with Salim Khazem on Machine Learning, and on how to prepare tasteful coffee, and also the time spent with Luis Felipe Wolf Batista and Devarsi Rawal. I enjoyed the time I had with Georges Chahine and Pete Schroepfer, during our field experiments with the crawler when the peculiar weather in Metz allowed it.

I am grateful to all the people from Cetim I had the opportunity to work with, and who made enjoyable the sessions of experiments we did together. I would like to thank Quang Vu and Henri Walaszek, in particular, for all their time and the friendly welcome they offered me, every time I came to Senlis.

I would like to thank all the members of the committee who kindly accepted to evaluate my thesis. Last, I would like to warmly express my gratitude to my family, my parents and my sisters, who have always been there, by my side.

*We know very little, and yet it is astonishing that we know so much,
and still more astonishing that so little knowledge can give us so much power.*

Bertrand Russell

To my family.

Abstract

The inspection of large plate-based metal structures, such as storage tanks or ship hulls, is a significant stake in the industry, which necessitates reliable and time-efficient solutions. Over the time, concealed defects, such as corrosion patches growing from the inner side of ship hulls, can develop, and detrimentally affect the structural integrity. To avoid catastrophic failures, the structures under operation must be inspected on a regular basis, so that structural maintenance can be scheduled when defects are detected. Current inspection methods traditionally rely on visual inspection, or on point-by-point thickness measurements, often performed by human operators. However, these methods are time-consuming, and cannot reasonably be used for an accurate and complete inspection of large metal surfaces.

This thesis considers the deployment of a mobile robotic system, equipped with an acoustic sensor for emission and reception of omnidirectional ultrasonic guided waves, with the aim to assess the integrity of large structures. Guided waves can propagate inside the structure material, over large distances, following a direction parallel with the surface. What make them appealing is that they are sensitive to the material integrity. In this context, the target objective is to enable long-range robotic inspection on large metal structures by listening to the wave reflections on structural features. Due to the possibility to localize potential defects at long distance all along the robot trajectory, long-range robotic inspection has the potential to be faster, and much more efficient than traditional inspection methods.

The abilities to accurately recover the robot position and to precisely map acoustic scatterers are critical for a successful inspection operation. For this reason, this thesis focuses on Simultaneous Localization and Mapping (SLAM) problems. More specifically, we propose various frameworks to recover the robot trajectory and map the boundaries of individual metal panels (which are the constitutive element of large metal structures) by relying on the ultrasonic echoes. Due to the wave omnidirectionality, the measurements contain information from all directions. This results in significant data complexity, and calls for appropriate signal processing methods. We will show that, by relying on ultrasonic wave reflections on the boundaries of a metal panel, the combination of a simple Delay-and-Sum (DAS) beamforming method with a wave propagation model can result in accurate localization and mapping results. The particularity of the proposed approach is that no explicit echo detection, nor echo association are required, whereas these are challenging tasks, especially in noisy environments. We propose algorithmic solutions to solve both the online SLAM and the full SLAM problems. We also tackle specific issues that arise when relying on DAS beamforming, such as signal interference, and the lack of uncertainty assessment on the map. An exploration strategy based on frontier exploration is also briefly presented, with the aim to enable the autonomous mapping of a metal surface. Eventually, we propose a method to automatically calibrate the wave propagation model when the propagation properties are not precisely known.

The proposed methods are validated by performing experiments in laboratory conditions, either with measurements that are collected by manually moving the transducers, or by using a real robotic platform. The results demonstrate the efficiency of our methods for on-plate localization and mapping purposes, and may support the emergence of a true long-range robotic inspection operation.

Résumé en français

1 L'inspection de larges structures en métal

Le transport efficace et fiable de marchandises autour du monde est un enjeu crucial pour le commerce mondial. Il est estimé qu'à tout moment, environ 50 000 navires, de taille moyenne ou large, parcourent les océans. Par ailleurs, chacun d'entre eux peut transporter jusqu'à plusieurs centaines de tonnes de marchandises. Dans ce contexte, afin de prévenir des accidents avec des conséquences potentiellement économiques, environnementales et humaines, l'inspection et la maintenance de la structure des navires sont des procédures critiques qui doivent être régulièrement effectuées. Au cours du temps, des défauts structurels, tels que des couches de corrosion, peuvent se développer. Comme ces derniers peuvent se former sur la partie intérieure des coques de navires, ils ne sont pas toujours directement visibles. Pour ces raisons, des procédures d'inspection fiables sont nécessaires pour l'industrie marine. De plus, étant donné que cette problématique concerne également d'autres structures en métal telles que les réservoirs de stockage, la détection de défauts non directement visibles est un problème industriel majeur.

Les méthodes d'inspection traditionnelles consistent principalement à l'inspection visuelle, et à la réalisation de mesures d'épaisseur point par point à l'aide d'une sonde ultrasonore, après avoir placé le bateau en cale sèche. Ces méthodes sont sous-optimales et lentes à réaliser, car le bateau peut être immobilisé pendant plusieurs jours, entraînant des pertes économiques substantielles pour son propriétaire. Par ailleurs, en raison de la taille limitée de la sonde acoustique utilisée pour les mesures d'épaisseur, et au vu de la taille importante de la coque, il est pour l'instant inenvisageable de réaliser une inspection qui soit à la fois complète et précise de l'ensemble de la surface, en un temps raisonnable.

L'inspection de structures sur de larges distances est un sujet de recherche actif. Les ondes ultrasoniques guidées, en particulier, sont de plus en plus déployées sur des systèmes modernes de contrôle non destructif. Ces ondes peuvent être émises à l'intérieur de matériaux en plaques par application de capteurs piézoélectriques. Elles peuvent se propager radialement autour de l'émetteur, sur de longues distances, et suivant une direction parallèle avec la surface. Le fait que ces ondes soient sensibles aux éléments structurels, tels que des discontinuités de la matière, rend leur utilisation attrayante. De façon générale, les ondes ultrasoniques guidées sont déployées sur des réseaux de capteurs statiques qui sont en permanence attachés à la structure pour le contrôle continu d'une zone restreinte. Cependant, dans l'industrie, ces ondes ne sont pas encore déployées sur des systèmes mobiles pour l'inspection de larges structures.

2 Localisation et cartographie en simultané par ondes guidées

L'inspection robotique longue distance de larges structures en métal, telles que les réservoirs de stockage ou les coques de navires, n'a pas encore été établie, alors que celle-ci peut être beaucoup plus rapide, précise, et moins sujette à des risques en comparaison avec les méthodes traditionnelles. Pour ce type d'application, un système robotique mobile serait équipé d'un capteur ultrasonique pour l'émission et la réception d'ondes ultrasonores. Suivant une trajectoire optimale pour l'inspection sur une large structure (tel qu'illustré sur la figure 1), les réponses ultrasonores de différentes zones pourrait être enregistrées. Dans cette configuration en pulse-écho, les mesures contiendraient les échos dus à des réflexions de l'onde excitée sur des éléments structurels, comme illustré par la figure 2. Après un traitement approprié des données obtenues, on pourrait obtenir une carte précise de la structure mettant en évidence les réflecteurs acoustiques,



Figure 1: Image d'un *crawler* magnétique navigant sur la façade d'un réservoir de stockage.

tels que les joints de soudures, les raidisseurs et les potentiels défauts.

Toutefois, la cartographie de réflecteurs acoustiques nécessite la connaissance de la position des points de mesures. Sur une large structure, la localisation précise du capteur embarqué (*i.e* avec une précision de l'ordre du centimètre) est essentielle pour la détection de défauts qui agissent comme des petits disperseurs acoustiques. Pour obtenir un tel niveau de précision, les données provenant de différents capteurs embarqués – par exemple une centrale inertielle, un capteur à bande ultra-large en plus du capteur ultrasonore – pourraient être *fusionnées*. En effet, en se basant sur la réflexion des ondes guidées sur les éléments structures tels que les joints de soudure, ces derniers pourraient être utilisés comme points de repère pour la localisation. L'intégration de mesures provenant d'un système de positionnement global (GPS), bien que courante pour la navigation en milieu extérieur non obstrué, n'est pas pertinente pour la localisation précise sur une large structure quasi verticale.

En somme, l'un des éléments primordiaux pour permettre de développement de l'inspection robotique longue distance est une solution au problème de localisation et cartographie en simultané (on utilisera l'acronyme anglais SLAM). En effet, à partir des mesures collectées le long de la trajectoire du robot, celle-ci pourrait être précisément reconstruite dans un premier temps, et une carte des éléments structurels servant à la localisation (tels que les joints de soudure) pourrait être déterminée. Dans une seconde étape, la trajectoire et la carte estimées seraient mises à profit pour inférer la position de défauts potentiels. Bien qu'une telle application apporterait




Figure 2: Une mesure ultrasonore obtenue expérimentalement par application de capteurs piézoélectriques sur une plaque en acier. Le signal contient de multiples paquets d'onde dûs aux réflexions de l'onde émise sur les bords de plaque.

des bénéfices majeurs à l'industrie, son développement est un problème de recherche significatif.

En comparaison avec les ondes directionnelles (qui peuvent seulement être émises et reçues pour une direction donnée), l'utilisation d'ondes omnidirectionnelles guidées peut mener à une inspection de l'environnement beaucoup plus rapide et efficace, car une seule mesure contient des informations provenant de toutes les directions. En revanche, cela requiert des méthodes de traitement du signal sophistiquées. Par exemple, pour localiser les bords d'une plaque en métal à partir de mesures ultrasonores (semblables à celle montrée par la figure 2), il est d'abord nécessaire d'extraire les échos du bruit, de regrouper entre eux les échos provenant d'une réflexion sur un même bord, et enfin estimer la position de ce dernier. Ces tâches difficiles doivent être réalisées pour des rapports signal-sur-bruit parfois très faibles, en présence d'échos d'ordre élevé qui ne sont pas labellisés, en présence d'incertitude sur la localisation du capteur, et sur la propagation du signal dans le guide d'onde, qui est un phénomène physique complexe. Dans le cas d'une tâche réelle d'inspection, il faudrait surmonter ces problèmes pour des distributions arbitraires de réflecteurs acoustiques.

Étant donné que de tels problèmes n'ont été que peu étudiés dans la littérature scientifique, cette thèse élabore différentes solutions pour résoudre le problème de localisation et cartographie en simultané en utilisant des ondes ultrasoniques omnidirectionnelles. Le but est d'évaluer leur potentiel pour une possible tâche d'inspection robotique sur une large structure. Une des contributions majeures de cette thèse est l'intégration de techniques de formation de voies (connue sous le nom de *beamforming* en anglais), qui est une technique classique de filtrage spatial, dans les approches Bayésiennes traditionnelles qui sont utilisées en robotique, tels que FastSLAM, le filtre Bayésien binaire, ou bien GraphSLAM. Les techniques de formation de voies permettent de sélectionner spatialement les signaux reçus dans le but de réaliser différentes tâches comme la localisation d'une source acoustique. Nous explorons comment ces techniques peuvent être employées pour résoudre une variété de problèmes de SLAM. Un des avantages de l'utilisation de la formation de voies, est que la détection explicite des échos et leur association ne sont pas nécessaires. Cela mène à des algorithmes de SLAM qui sont plus simples, et plus robustes au bruit. De plus, alors que la formation de voies est généralement utilisée pour des données acquises par des chaînes de capteurs, nous verrons qu'elle peut également être appliquée à des données acquises par un seul capteur mobile, lequel créant alors une chaîne "synthétique" de capteurs. En raison de la complexité de l'application visée, cette thèse se focalise principalement sur la cartographie des bords d'une plaque en métal individuelle, élément constitutif des structures larges en métal, et n'aborde pas directement la détection de défauts. Toutefois, comme les joints de soudure réfléchissent vraisemblablement l'énergie acoustique de façon plus importante que des défauts sur une large structure réelle, il est considéré que leur identification est une étape essentielle avant que la détection de défauts puisse être faisable.

3 Plan de la thèse

Dans cette thèse, nous proposons des solutions à différents problèmes de localisation et de reconstruction de géométrie de plaques par ondes guidées ultrasonores. Initialement, les méthodes présentées reposent sur de la connaissance à priori forte, ainsi que sur des hypothèses sur la géométrie de la plaque. Par la suite, ces hypothèses sont relaxées afin d'étendre les méthodes proposées à des cas plus généraux. Le contenu de cette thèse est comme suit.

Le **Chapitre 2** introduit les différentes notions théoriques qui sont nécessaires pour la compréhension des méthodes développées dans cette thèse. En particulier, les éléments théoriques fondamentaux sur la propagation des ondes guidées sont rappelées, la notion de filtrage spatial

par formation de voies est présentée, et les méthodes conventionnelles de filtrage Bayésien pour la localisation et la cartographie sont introduites.

Dans le **Chapitre 3**, nous présentons un simple filtre à particules dont le modèle d'observation est obtenu par corrélation des signaux ultrasoniques avec un modèle de propagation de l'onde, dans le but de permettre la localisation précise du capteur sur une plaque de métal rectangulaire. La particularité de cette approche réside dans le fait que ni la détection, ni l'association explicites des échos ne sont nécessaires. L'efficacité de cette approche est évaluée sur des mesures ultrasoniques obtenues par déplacement manuel des capteurs acoustiques sur une plaque en métal, dans un environnement de laboratoire. Les résultats montrent que dans la configuration considérée, une précision centimétrique sur la position des capteurs peut être obtenue.

Dans le **Chapitre 4**, nous proposons un algorithme FastSLAM qui intègre la technique de formation de voies par délai et somme pour résoudre le problème de SLAM dit "en ligne" (ou en temps réel), ce qui permet de retrouver la position des bords d'une plaque rectangulaire de façon conjointe avec la localisation des capteurs. Les performances de cette méthode sont comparées avec celles d'une approche de FastSLAM qui se base sur la détection et l'association explicites des échos. Les résultats montrent que, avec la nouvelle méthode proposée, la localisation et la cartographie se fait de façon beaucoup plus précise et robuste. Par ailleurs, une expérience de SLAM est réalisée avec un *crawler* magnétique réel dans le but de démontrer l'applicabilité pratique du SLAM ultrasonique.

Le **Chapitre 5** présente une approche de cartographie plus élaborée que celle introduite dans le Chapitre 4, car elle permet de construire une représentation probabiliste en grille de la surface d'une plaque de métal, conjointement avec la construction d'une représentation par éléments caractéristiques du même environnement. L'objectif de cette approche est de permettre l'évaluation de l'incertitude sur la reconstruction de la carte, et de mettre à profit la contrainte de consistance entre les deux représentations pour mitiger l'effet d'interférence du signal induit par la formation de voies. Cette nouvelle approche de cartographie est, par ailleurs, intégrée dans une stratégie d'exploration pour conduire, de façon automatique, un agent autonome vers les positions d'acquisitions qui sont les plus informatives. Enfin, notre approche de cartographie par représentations combinées est utilisée pour étendre notre approche de reconstruction géométrie à des polygones plus arbitraires que des rectangles.

Dans le **Chapitre 6**, nous proposons une approche de GraphSLAM pour résoudre le problème de SLAM dit "complet" en utilisant également les échos d'ordre supérieur qui sont contenus dans les mesures ultrasonores. Cette approche intègre un modèle d'observation d'échos d'ordre supérieurs, ainsi qu'une technique de formation de voies dans un optimiseur de moindres carrés non-linéaire. Cette approche est validée sur des données expérimentales.

Le **Chapitre 7** présente une approche de formation de voies optimale, de sorte à rendre possible la cartographie d'une structure en métal sans connaissance précise à priori du modèle de propagation des ondes. En effet, celle-ci risque de manquer lors d'une tâche d'inspection réelle. Nous montrons que, par la maximisation du niveau de focalisation de la technique formation de voies, les paramètres du modèle de propagation peuvent être adaptés afin d'obtenir des résultats précis de cartographie, et lorsque les propriétés de propagations ne sont pas parfaitement connues. Pour mettre en évidence le niveau d'efficacité de cette approche, celle-ci est évaluée sur des mesures obtenues dans des conditions d'acquisition nominales, et non nominales.

Enfin, le **Chapitre 8** résume l'ensemble des contributions de cette thèse et propose au lecteur des suggestions de travail futur.

La table suivante présente un résumé des différents problèmes de localisation et de cartographie qui sont considérés dans cette thèse, ainsi que des hypothèses qui sont faites.

Chapitre	Problème	En ligne / Hors ligne	Géométrie de plaque	Modèle de propagation
3	Localisation	En ligne	Rectangulaire	Connu
4	SLAM	En ligne	Rectangulaire	Connu
5.1	Cartographie	En ligne	Rectangulaire	Connu
5.2	Exploration	En ligne	Rectangulaire	Connu
5.3	Cartographie	Hors ligne	Polygonale	Connu
6	SLAM complet	Hors ligne	Rectangulaire	Connu
7	Cartographie	Hors ligne	Rectangulaire	Inconnu

Table 1: Résumé des différents problèmes considérés dans cette thèse selon chaque chapitre.

Contents

Chapter 1 Introduction	1
1.1 On the inspection of large metal structures	1
1.2 Acoustic SLAM using omnidirectional guided waves	2
1.3 Thesis outline	4
1.4 List of publications	5
Chapter 2 Background theory	7
2.1 Lamb wave propagation in a metal plate	7
2.1.1 Rayleigh-Lamb equations	8
2.1.2 Green's function	9
2.2 Array signal processing	10
2.2.1 Introduction to beamforming	10
2.2.2 Applications in Structural Health Monitoring	11
2.3 Bayesian filtering	12
2.3.1 Probabilistic interaction models	12
2.3.2 Recursive Bayes filter	14
2.3.3 Particle filter	15
2.3.4 Mapping	16
2.3.5 Simultaneous Localization and Mapping	17
Chapter 3 On-plate Monte-Carlo localization	19
3.1 Introduction	19
3.1.1 Related works	20
3.2 On-plate localization using ultrasonic guided waves	21
3.2.1 Problem statement	21
3.2.2 Measurement model and image sources	21
3.2.3 Correlation signal	23
3.2.4 Monte-Carlo localization	24
3.3 Experimental results	25

3.3.1	Experimental setup	25
3.3.2	Dispersion curves	26
3.3.3	Echo detection	27
3.3.4	Localization results	28
3.4	Chapter summary	31
Chapter 4 Online Simultaneous Localization and Mapping		33
4.1	Introduction	33
4.1.1	Acoustic mapping	34
4.1.2	Simultaneous Localization and Mapping	34
4.1.3	Introduction to FastSLAM	35
4.2	Proof-of-concept FastSLAM method	36
4.2.1	Notations and assumptions	38
4.2.2	Echo detection	38
4.2.3	From echoes to edge hypotheses	39
4.2.4	FastSLAM integration	40
4.2.5	Overview on the SLAM results	41
4.3	A FastSLAM approach integrating beamforming maps	43
4.3.1	Plate boundary mapping with delay-and-sum beamforming	44
4.3.2	Particle evaluation and FastSLAM algorithm	46
4.4	Experimental results	47
4.4.1	Experimental setups	47
4.4.2	Localization and mapping results	48
4.4.3	Experiment with a magnetic crawler	52
4.5	Chapter summary	55
Chapter 5 Guided wave-based mapping with combined representations		57
5.1	Motivation and background	57
5.1.1	Related works	59
5.1.2	Binary Bayesian filtering	59
5.2	Grid and feature-based mapping using ultrasonic guided waves	61
5.2.1	Feature-based mapping via beamforming	61
5.2.2	Estimation problem formulation	62
5.2.3	Binary filtering with adaptive inverse sensor models	63
5.2.4	Inverse model estimation in the Cartesian space	64
5.2.5	Inverse model estimation in the line boundary space	66
5.2.6	Algorithm implementation	67

5.3	Numerical and experimental results	67
5.3.1	Results with hand-picked measurements	67
5.3.2	Simulation results	69
5.3.3	Results with a magnetic crawler	69
5.4	Towards ultrasound-based autonomous exploration	71
5.4.1	Related works	71
5.4.2	Overview of the frontier exploration approach	72
5.4.3	Overview of the exploration results	73
5.5	Polygonal shapes reconstruction from acoustic echoes	76
5.5.1	Related works	77
5.5.2	Problem statement	77
5.5.3	A regularized cost for polygonal shape reconstruction	78
5.5.4	Shape reconstruction with simulated annealing	79
5.5.5	Simulation results	80
5.5.6	Experimental results	82
5.6	Chapter summary	84
Chapter 6	Graph-based Simultaneous Localization and Mapping	85
6.1	Motivation and background	85
6.1.1	Multi-order echoes signal processing	86
6.1.2	Related works on GraphSLAM	86
6.2	A graph-based framework for ultrasonic SLAM	87
6.2.1	Localization with a multi-order echo observation model	87
6.2.2	Mapping only with beamforming	88
6.2.3	The GraphSLAM cost criterion	90
6.2.4	Sparse non-linear least squares minimization	91
6.3	Experimental results	92
6.3.1	Comparison of the observation models performance	93
6.3.2	Experimental localization and mapping results	95
6.3.3	A study on the sensitivity of GraphSLAM to the loss parameter	99
6.3.4	Discussions	100
6.4	Chapter summary	101
Chapter 7	Learning the propagation model for Lamb wave-based mapping	103
7.1	Introduction and background	103
7.1.1	Related works	104
7.2	Proposed method	105

7.2.1	Lamb wave-based mapping	105
7.2.2	Optimal beamforming for model learning	106
7.2.3	Optimization with simulated annealing	108
7.3	Experimental results	109
7.3.1	Experimental scenarios	109
7.3.2	Influence of the high-pass filter	110
7.3.3	Correlation between the loss and the mapping accuracy	111
7.3.4	Mapping results with model learning in nominal conditions	113
7.3.5	Mapping results with model learning in non-nominal conditions	114
7.3.6	Evaluation of our approach with sparse measurements	115
7.3.7	Discussions	117
7.4	Chapter summary	118
Chapter 8 Conclusion		121
8.1	Summary of the contributions	121
8.2	Perspectives	122
Appendix: Derivation of the integrals of the inverse models		123
A:	Derivation of the first integral	123
B:	Derivation of the second integral	123
Bibliography		125

List of Figures

1	Image d'un <i>crawler</i> magnétique navigant sur la façade d'un réservoir de stockage.	vii
2	Une mesure ultrasonore obtenue expérimentalement par application de capteurs piézoélectriques sur une plaque en acier. Le signal contient de multiples paquets d'onde dûs aux réflexions de l'onde émise sur les bords de plaque.	vii
1.1	(Left) Closeup of the Altiscan magnetic crawler on a metal structure. (Right) Ultrasonic guided waves reflecting on the edges of a metal panel in a simulation environment.	1
1.2	The Altiscan crawler navigating on the facade of a 20 m high storage tank. . . .	2
1.3	An ultrasonic measurement acquired experimentally by applying contact piezoelectric transducers on a steel plate sample. The measurement contains multiple wave packets due to the reflections of the incident wave on the plate boundaries.	2
2.1	Simplified representation of the out-of-plate distortion induced by Lamb wave symmetric and anti-symmetric modes in a plate. The plate boundaries at equilibrium state of the plate are represented by the discontinuous lines.	7
2.2	Illustration of the wave dispersion effect in a metal plate with simulated data. . .	8
2.3	Dispersion curves for several low order symmetric and anti-symmetric Lamb wave modes in a aluminum plate ($c_L = 6420$ m/s, $c_T = 3040$ m/s). The top figure depicts the group velocities while the bottom plot shows the wavenumbers.	9
2.4	Illustration of delay-and-sum beamforming to achieve source localization. The filter weights $\mathbf{w}(\boldsymbol{\theta})$ apply delays to the input signals based on the distance between each antenna element and candidate position $\boldsymbol{\theta}$. The beamformer outputs the sum of the delayed signals which presents constructive interference when the candidate $\boldsymbol{\theta}$ matches the true source location.	11
2.5	Defect detection using a network of sensors. Baseline subtraction and a multi-channel processing technique are used to obtain imaging results. The illustrations are taken from [37].	12
2.6	Illustration of the robot state along with adopted notations in a 2D world.	13
2.7	Representation of the dynamic process, made of successive robot motions and perception actions, with a Bayesian network. The robot poses \mathbf{x}_i are related to the past one \mathbf{x}_{i-1} and odometry input \mathbf{u}_{i-1} , while measurements \mathbf{z}_i depend only on the current state \mathbf{x}_i	14
2.8	Representation of online SLAM with a Bayesian network. The robot poses \mathbf{x}_i are related to the previous state \mathbf{x}_i and odometry input \mathbf{u}_{i-1} , while measurements \mathbf{z}_i depend both on the current state \mathbf{x}_i and the map \mathbf{M}	17

3.1	Illustration representing a robotic agent whose pose is $\mathbf{x} = [x, y, \theta]^T$ on a rectangular metal plate of dimensions $w \times h$	22
3.2	Illustration of the image source model. On the left is an actual setup with an omnidirectional emitter/receiver situated in a space partially bounded by two reflective lines. On the right is the equivalent setup provided by the image source model. The first-order image sources at \mathbf{p}_1 and \mathbf{p}_2 are obtained by symmetry of the real source on each boundary. The second-order one at \mathbf{p}_3 is obtained by symmetry of one of the first-order image sources.	23
3.3	(Left) Picture of the experimental setup for data acquisition. (Right) Position of the measurement points on the plate.	26
3.4	Two-dimensional space-time Fourier transform of the experimental wave field recorded on the considered aluminum plate sample. It is overlaid by the theoretical dispersion curves of the A0 and S0 Lamb wave modes.	27
3.5	Illustration showing the position of the sensors on the metal plate for acquisition of the test measurement.	28
3.6	The acoustic measurement acquired on the aluminum plate after placing the transducers on the plate, as depicted in Fig. 3.5 (top plot), and the correlation signal obtained from the measurement along with its envelope (bottom plot). The ranges relative to the first-order reflections (8, 37, 52 cm) can be successfully retrieved from the local maxima of the envelope signal. The echo at nearly 45 cm corresponds to a higher-order reflection.	28
3.7	Qualitative localization results of the particle filter during measurement steps 1, 7, 14, 45, 83 and 108 of the simulated lawn mowing trajectory. The red dot represents the ground truth sensor position, the red line accounts for the true sensor trajectory. The blue rectangle represents the outline of the plate, and the arrows account for the particles' distribution.	29
3.8	Evolution of the localization errors (determined based on the median of the particles' state) on the x (left) and y (right) coordinates for 10 repetitions of the considered scenario. The red line is the average error calculated over the repetitions.	30
3.9	Evolution of the localization errors (determined based on the median of the particles' state) on the x (left) and y (right) coordinates for 10 repetitions of the considered scenario, and using the particle disturbance probability $\gamma = 0$. The red line is the average error calculated over the repetitions.	30
4.1	Line hypotheses generated from two measurements at position P_{t-1} and P_t . Left: the robot moves away from the edge and generate a single hypothesis, right: generic case leading to two edge hypotheses.	39
4.2	Illustration of the echo detection process. Panel a) shows the emitted signal. Panel b) shows the echo signal (blue trace) and the reconstructed impulse response (orange trace). Panel c) shows the reconstruction residual.	42
4.3	Evolution of the map representation for the highest-ranked particle at step 1, 36, 124 and 239. The purple frame represents the outline of the true plate. The red dot is the current estimated sensor pose, and the red line is the history of its estimated trajectory. The darkness of the line is proportional to the number of times they have been observed.	42
4.4	Representation of lines using (r, α) coordinates. As a convention, the reference frame is taken as the initial robot position, and the x axis is aligned with the initial robot heading.	45

4.5	Acquisition positions on the steel plate.	47
4.6	Trajectories estimated by all the particles (black lines), dead-reckoning trajectories (dash magenta lines) and map retrieved by the most likely particle (green lines) during Steps 1, 22, 50 and 108 for a lawn-mower path on plate 1 (zoom for details). The true outline of the plate and true sensor positions correspond to the blue rectangle and blue dot respectively.	48
4.7	Beamforming map for the particle with the highest weight during the final step. The (r, α) coordinates are expressed in the reference frame which origin is the initial position of the sensors, and the x axis is aligned with the initial heading. The rectangles indicate the edges retrieved with our method.	49
4.8	Localization and mapping results over 100 repetitions of a lawn-mower path on plate 1 for the previous and the new FastSLAM methods. a) Average estimation errors on the range parameter of the lines. b) Average estimation errors on the angle parameter. c) Average localization errors in the estimated plate frame. The 10% and 90% quantiles correspond to the upper and lower bounds of the coloured areas. The scales along the y-axis are logarithmic.	50
4.9	Localization and mapping results over 100 repetitions of a lawn-mower path on plate 2 for the new FastSLAM method. a) Average estimation errors on the range parameter of the lines. b) Average estimation error on the angle parameter. c) Average localization errors in the estimated plate frame. The 10% and 90% quantiles correspond to the upper and lower bounds of the colored areas. The scales along the y-axis are logarithmic.	51
4.10	Schematic of designed electrical circuit for a true pulse-echo setup. The diodes in parallel are 1N4148 commutation diodes, while the diodes in series are Zener diodes.	53
4.11	Qualitative SLAM results obtained with a magnetic crawler, and using the proposed FastSLAM approach, during different steps of the experiment.	54
5.1	An example of an occupancy grid with coarse resolution. It illustrates the typical effect of integrating a single range measurement to an obstacle into an initial uncertain map.	60
5.2	Illustration of the interference effect. (a) shows mapping results obtained in a simulated scenario, where one edge is not correctly identified from the beamforming map shown in (b). The origin O of the reference frame is indicated along with the (r, α) coordinate of the left line. The ground truth plate outline is represented by the blue rectangle.	62
5.3	Illustration of beamforming map filtering with the first acoustic echo for interference rejection. a) shows the mapping results obtained after applying, to the beamforming map in Fig. 5.2b, a binary mask that filters lines which can no longer exist, as they cross the area identified as "on the plate". The on-plate area identified with the ground-truth echo measurements is represented in (a) as the lighter area. The origin O of the reference frame is indicated along with the (r, α) coordinate of the left line.	62
5.4	Illustration of the inverse model in the Cartesian space for different values of detected range (which is equal to the radius of the light disk surface). The sensor position is represented by the red point, and the considered line estimate is in green.	66

5.5	Mapping results achieved on the aluminum plate during Steps 1, 42, 108 and 216. Each sub-figure depicts, on the left, the sensor trajectory along with the detected range to the closest edge represented by the circles, and on the right, the filtered beamforming maps during the same step, where the red rectangles indicate the retrieved edge estimates.	68
5.6	Final mapping errors w.r.t. the noise level on the echo detection assessed over 10 repetitions of the lawnmower trajectory. The errors are determined both for our approach and for the baseline one. The left plot shows the range errors, while the orientation errors are shown on the right plot. The solid lines represent the average values, while the upper and lower bounds of the colored areas are situated respectively above and below one standard deviation to the mean values.	69
5.7	Combined feature-based and grid-based mapping with a robotic platform relying on UGWs. (a) shows the experimental setup. (b), (c), (d) and (e) present the mapping results at different steps along the robot transect. The robot pose (red arrow) and trajectory (red line) are also represented.	70
5.8	Illustration of a mapping output provided by our approach. The blue dot represents the sensor position and the blue line is its overall trajectory. The red rectangle represents the estimate $\hat{\mathbf{M}}_t$ obtained via DAS beamforming. The grid $\hat{\mathcal{O}}_t$ is accounted for with the color of the cells. The valid (resp. invalid) frontier points corresponds to the red (resp. black) dots. The green point indicates the selected next best point.	73
5.9	Mapping outputs obtained in simulation during different steps of the mapping process. The blue dot represents the sensor position, and the blue line is the sensor trajectory. The red rectangle represents the feature-based map estimate obtained via DAS beamforming. The grid $\hat{\mathcal{O}}_t$ is accounted for with the color of the cells. The valid (resp. invalid) frontier points correspond to the red (resp. black) dots. The green point indicates the selected next best point.	74
5.10	Evolution of the plate surface coverage for the four exploration strategies with respect to the traveled distance in the simulations. The discontinuous lines represent the coverage mean value over 50 repetitions of each method, while the upper and lower bounds of the colored areas represent respectively the 90% and 10% quantiles.	75
5.11	Top view of the experimental setup. A turtlebot platform is used along with AR tag markers for ground truth positioning.	75
5.12	Evolution of the plate surface coverage for the three exploration strategies with respect to the traveled distance during the experiments. The solid lines represent the coverage mean value over 5 repetitions of each method, while the upper and lower bounds of the colored areas represent respectively the 90% and 10% quantiles.	76
5.13	Schematic of the problem setup and notations. An example of shape to estimate is represented by the red polygon. A candidate geometry estimate is represented by the set of green lines, whose (r, α) coordinates are only provided for one of them. The area A formed with a single range measurement ρ is represented by the lighter area. The inner surfaces of A and $\hat{\mathbf{M}}$ are delimited with the discontinuous dashed lines.	78

5.14	The different simulated setups along with an instance of mapping results. On the left of each sub-figure, the outline of the ground-truth geometry is provided in discontinuous red, the green lines account for the estimated edges, and the robot path and acquisition positions are represented in black. The corresponding beamforming maps are provided on the right, where the red rectangles indicate the retrieved edges. The edge coordinates are expressed in the fixed frame, whose origin is the initial sensor position O , and the x -axis is aligned with the initial robot heading.	81
5.15	Evolution of the performance criteria for the 3 scenarios, and for 10 repetitions of the optimizer. The top plots show, in blue, the evolution of the costs. The middle plots depict the estimated number of edges after application of the edge rejection criterion. The bottom plots show the IoU between estimated and ground-truth surfaces. The median values are in red.	82
5.16	Mapping results for the experimental setup. (a, left) shows the simulated sensor path (in black), the plate outline (discontinuous red) and estimated edges (green). The corresponding beamforming map is provided on the right, where the red rectangles indicate the retrieved edges. (b) shows the evolution of the loss, the number of estimated edges and the value of $\text{IoU}(\hat{\mathbf{M}}, \mathbf{M})$ for 10 repetitions of the optimizer. The black horizontal line on the IoU plot is the value achieved with the baseline method. The median values are in red.	83
6.1	Positions of the image sources (red dots) for a rectangular plate with dimensions $w \times h$ (blue), and a fixed emitter position $[x, y]$ (blue dot).	88
6.2	Illustration of the envelope signal obtained from an experimental ultrasonic measurement acquired on a rectangular metal panel, and comparison with the expected envelope calculated with the matched Gaussian mixture model.	89
6.3	Graph representation of the (sparse) SLAM problem, where the variables to estimate are the positions \mathbf{x}_i and the map \mathbf{M} , the inputs are the odometry data \mathbf{u}_i and measurements \mathbf{z}_i . The constraints (induced by the external measurements) between the variables are represented by the graph edges, which are related to the l functions.	90
6.4	Graph representation of the SLAM problem after integration of DAS beamforming. The variables to estimate are the positions \mathbf{x}_i , and the map \mathbf{M} is hidden as it can directly be inferred from all \mathbf{x}_i using DAS beamforming. For simplicity of the graph representation, an edge is drawn to account for the existence of a constraint between two variables.	91
6.5	Likelihood maps of sensor presence obtained using measurements acquired on the aluminum plate (a), (b), and on the steel plate (c), (d). The true sensor position is indicated by the red mark. The outline of each plot accounts for the plate boundaries, and the color maps only spans the range from the minimum to the maximum likelihood value.	93
6.6	Average localization errors calculated using the two considered models. The position error is averaged over all the measurements acquired on the aluminium plate (top) and steel plate (bottom) with respect to γ , yielding the solid lines. The upper and lower bounds of the colored areas are situated respectively above and below one standard deviation from the average.	94

6.7	Ground truth trajectories and dead-reckoning trajectories used for Scenario 1 (a) and Scenario 2 (b). The black rectangle represents the actual outline of the metal plate.	95
6.8	Localization and mapping results obtained with FastSLAM for the Scenario 1. The figure represents the trajectories estimated by all the particles (red lines), along with the map retrieved by the particle with the highest weight (green lines) during Steps 1, 6, 18, 49, 88 and 108 of the simulated lawn-mowing trajectory. The true outline of the plate and true sensor positions correspond to the black rectangle and blue dot, respectively.	96
6.9	Representation of the SLAM results during iteration Steps 1, 15, 29, and 41 of the NLS optimization process for Scenario 1. The true outline of the plate is represented by the black rectangle. For better comparison of the reconstructed trajectory with the ground truth, they are expressed in the estimated (resp. ground truth) plate frame, while the origin of the estimated plate geometry is fixed at $[0, 0]$	97
6.10	Evolution of the loss (left), along with the localization and mapping errors (right) during the NLS optimization process for Scenario 1. The loss evaluated at ground truth is represented as a horizontal line.	98
6.11	Evolution of the loss (left), along with the localization and mapping errors (right) during the NLS optimization process for Scenario 2. The loss evaluated at ground truth is represented as a horizontal line.	98
6.12	Representation of the SLAM results during iteration step of the NLS optimization process for Scenario 2. The true outline of the plate is represented by the black rectangle. For better comparison of the reconstructed trajectory with the ground truth, they are expressed in the estimated (resp. ground truth) plate frame, while the origin of the estimated plate geometry is fixed at $[0, 0]$	99
6.13	Evolution of loss along with the localization and mapping errors during the NLS optimization in Scenario 1, and for different cost parameters λ . The top left graph shows the evolution of the difference between the achieved loss and the loss at ground truth. A close view of the same graph is provided on the bottom left plot. The middle plot shows the evolution of the localization errors, while the mapping errors are provided on the right plot.	100
7.1	Picture showing the experimental setup for Scenario 3.	110
7.2	Different beamforming maps computed using the data from Scenario 2 and the retrieved edges. a) shows the standard beamforming map, which yields correct estimates. b) depicts the map obtained with the same data sub-sampled in time. One edge is not correctly estimated. c) shows the high-pass filtered beamforming map obtained from the sub-sampled data. The correct plate geometry is recovered, while the fuzzy areas due to interference and higher-order reflections have been partially filtered out.	111
7.3	2D plots showing the correlation and loss values achieved using each of the propagation models from the predefined model set, and for the three scenarios. On the left, the loss value and geometry estimates are determined from the high-pass filtered beamforming maps. On the right, they are determined from the standard maps without filtering.	112

7.4	Evolution of the loss value and reconstruction errors for 10 repetitions of the optimization process with simulated annealing, and using the data from Scenario 1 (a) and Scenario 2 (b). The solid lines represent the mean values. The upper and lower bounds of the filled areas represent the minimum and maximum values respectively during each iteration. For comparison, the values achieved with the predetermined propagation model for each scenario are displayed as horizontal lines. The scale along the y-axis is logarithmic for the range error plots only. . . .	114
7.5	Evolution of the loss value and reconstruction errors for 10 repetitions of the optimization process and using the data from Scenario 3. The blue lines represent the mean values. The upper and lower bounds of the colored areas represent the minimum and maximum values respectively during each iteration. For comparison, the values achieved with the predetermined propagation model are displayed as horizontal red lines. The scale along the y-axis is logarithmic for the range error plot only.	115
7.6	Simulated robot paths used for the experiments on the steel plate (left) and the aluminum plate (right). The path lengths (<i>i.e</i> the number of acquisition positions) are respectively 36 and 26.	116
7.7	Reconstruction errors w.r.t. the number of measurements considered in the loss, and using the data from Scenario 2 (a) or from Scenario 3 (b). The measurements are integrated one after the other in the loss, following the simulated paths. The error is only evaluated for the model yielding the lower loss value for the two cases (model selection with and without filtering the beamforming maps). For comparison, the errors achieved with the predetermined models are also displayed.	116
7.8	Loss value computed on the high-pass filtered beamforming maps with respect to the longitudinal and transverse velocities. The data from Scenario 1 are used for the plot on the left, while the data from Scenario 2 are used for the plot on the right.	117

Acronyms

DAS : Delay And Sum
DoA : Direction of Arrival
GPS : Global Positioning System
GWs : Guided Waves
IMU : Inertial Measurement Unit
ML : Maximum Likelihood
NDE : Non-Destructive Evaluation
NDT : Non-Destructive Testing
NLS : Non-linear Least Squares
OG : Occupancy Grid
SA : Simulated Annealing
SH : Shear Horizontal
SHM : Structural Health Monitoring
SLAM : Simultaneous Localization and Mapping
SNR : Signal to Noise Ratio
UGWs : Ultrasonic Guided Waves
UWB : Ultra Wide Band

Chapter 1

Introduction

1.1 On the inspection of large metal structures

The efficient and reliable shipping of merchandise across the world is a cornerstone of world trade. It is estimated that, at any time, there are around 50 000 medium and large merchant ships crossing the oceans, each of them carrying up to hundreds of thousands of tons of goods [1]. In this context, to prevent catastrophic failures with large economical, environmental and life-threatening consequences, the assessment and maintenance of marine vessels' structural integrity is a critical stake, and it is to be carried out on a regular basis. Over the time, structural damages, like corrosion patches that locally affect the thickness of the ship hull, may develop. As they can appear from the inner side of the hull, they may not directly be visible. For these reasons, reliable procedures for assessing ships' hull structural integrity are needed in the marine industry. As it may also concern other large metal structures such as storage tanks, the detection of concealed defects is a significant issue in industry.

Current inspection methods essentially rely on visual inspection and point-by-point thickness measurements (either performed by human operators or by a remote-controlled magnetic crawler, as in Fig. 1.1), when the ship is in dry docks. These methods are suboptimal and time-consuming, as the ship can be immobilized for up to 8 days, with significant financial impacts for shipowners [1]. Besides, due to the limited size of the probes used for the thickness measurements, and based on the large size of the structures of interest, one cannot consider inspecting the entire surface within a reasonable amount of time.




Figure 1.1: (Left) Closeup of the Altiscan magnetic crawler on a metal structure. (Right) Ultrasonic guided waves reflecting on the edges of a metal panel in a simulation environment.

The inspection of structures over long ranges is being actively investigated. Ultrasonic Guided Waves (UGWs), in particular, are being integrated into modern Non-Destructive Testing (NDT)-capable devices. These waves can be emitted in plate materials using contact piezoelectric transducers, and can propagate radially over long distances, in a direction parallel to the surface. What makes them appealing is that they are sensitive to structural features, such as a material discontinuity (illustrated in Fig 1.1), or defects. Usually, they are deployed on static networks of sensors that are permanently attached to the structure to monitor local areas. However, in the industry, they are not yet deployed on mobile systems for NDT on large structures.



Figure 1.2: The Altiscan crawler navigating on the facade of a 20 m high storage tank.

1.2 Acoustic SLAM using omnidirectional guided waves

The long-range robotic inspection of large metal structures, such as storage tanks or ship hulls, has not yet been established, whereas it has the potential to be faster, more accurate, and less prone to risks than traditional methods. In this envisioned application, a robotic platform would be equipped with an ultrasonic sensor for both signal excitation and reception. Following an optimal trajectory on a large structure, as illustrated in Fig. 1.2, the ultrasonic response of various areas of the overall structure can be collected. In this pulse-echo setup, the measurements would contain ultrasonic echoes due to reflections of the excitation wave on structural features, as illustrated in Fig 1.3. After applying appropriate signal processing to the ultrasonic measure-




Figure 1.3: An ultrasonic measurement acquired experimentally by applying contact piezoelectric transducers on a steel plate sample. The measurement contains multiple wave packets due to the reflections of the incident wave on the plate boundaries.

ments, one could obtain an accurate map of the structure, depicting structural reflectors such as weld joints, stiffeners, and potential defects. The development of long-range robotic inspection is considered by the European project BugWright2 [1], whose main objective is to enable the inspection of ship hulls by a team of collaborative robots.

It is acknowledged that mapping acoustic scatterers necessitates prior knowledge of the measurement positions. On a large metal structure, accurate positioning (*i.e* probably at centimeter precision) would be critical to detect tiny acoustic reflectors, such as defects. To obtain such level of precision, one may *fuse* the information provided by different sensors embedded on the robotic system: namely an Inertial Measurement Unit (IMU), an Ultra Wide Band (UWB) sensor, along with an ultrasonic transducer. Indeed, by relying on ultrasonic echoes due to reflections on structural features such as weld joints, one may use the latter as structure-bound landmarks for localization. The integration of Global Positioning System (GPS) measurements, though widely used for navigation in outdoor and clutter-free environments, may not be relevant for accurate localization on a nearly vertical facade.

Overall, one of the key elements necessary to enable long-range robotic inspection is a solution to the ultrasonic Simultaneous Localization and Mapping (SLAM) problem. Indeed, based on the collected measurements, one would first attempt to accurately recover, with the aid of other external localization systems, the robot trajectory along with the position of structural features that act as acoustic reflectors (but are not defects). In a second step, one would use the estimated trajectory and map to infer the position of potential defects. Although such an application would bring major benefits to industry, it is still a significant research issue.

Compared to directional waves (that can only be emitted and received in a given direction), the use of omnidirectional guided waves can lead to a much more efficient probing of the environment, as one single measurement contain information from all the directions. However, this requires sophisticated signal processing. For example, to map the boundaries of a metal panel from ultrasonic measurements as depicted by Fig. 1.3, one may need to identify the actual echoes from the noise, to group together echoes that originate from a reflection on a same boundary, and to subsequently estimate the boundary location. These challenging tasks should be performed for low Signal-to-Noise Ratio (SNR) levels, in the presence of (unlabeled) multi-order reflections, localization uncertainty, and uncertainty on signal propagation in the wave guide, which is a complex physical phenomenon. During a realistic scenario, one may have to address these challenges for more arbitrary distributions of acoustic reflectors.

As such a research problem has been little investigated in the literature, this thesis elaborates various frameworks to solve the SLAM problem by relying on omnidirectional waves, with the aim to assess their potential for a long-range robotic operation, on a large metal structure. One of the core contributions is the integration of beamforming in classical estimation frameworks used in robotics, such as FastSLAM, the Bayesian binary filter, or GraphSLAM. Beamforming is a conventional technique in signal processing that can apply spatial selection to the received signals to achieve various tasks, such as the localization of a sound source. We explore how this technique can be leveraged to design solutions to a class of ultrasonic localization and mapping problems. One of the key advantages of this approach is that explicit echo detection, and echo association are not required, which leads to SLAM algorithms that have low algorithmic complexity, and that are robust to noise. Furthermore, while beamforming is often used in sensor arrays, we will see that this technique can seamlessly be applied to measurements acquired by a single robot which, by moving, creates a "synthetic" array. Due to the complexity of the target application, this thesis only focuses on recovering the boundaries of individual metal panels, thus not considering the detection of defects. Still, as weld joints may reflect the acoustic energy more strongly than defects, it is envisioned that their identification is an essential step prior to defect detection.

1.3 Thesis outline

In this thesis, we address various ultrasound-based localization and plate-geometry reconstruction problems. We initially rely on strong hypotheses, and prior knowledge on the plate environment to elaborate simple and sound solutions to the localization and mapping problems, while the subsequent chapters present how these hypotheses can be relaxed to extend the proposed frameworks to more general situations. The outline of this thesis is as follows:

Chapter 2 presents the background theory necessary to comprehend the localization and mapping techniques presented in this thesis. It recalls the fundamentals of wave propagation theory in a metal plate, introduces the notion of spatial filtering with beamforming, and briefly recalls the conventional frameworks for localization and mapping used in robotics.

Chapter 3 introduces a simple particle filter whose observation model is obtained by correlating the ultrasonic measurements with a wave propagation model, with the aim to achieve accurate robot localization on a rectangular metal panel by relying on UGWs. The particularity of this approach is that no explicit echo detection nor echo association is needed. The performance of this approach is evaluated by using measurements acquired by moving by hand piezoelectric sensors on a metal plate sample, in a laboratory environment.

In **Chapter 4**, we propose a FastSLAM algorithm that integrates Delay-and-Sum (DAS) beamforming to solve the online SLAM problem, which can additionally recover the geometry of a rectangular panel. The performance of this approach is compared with that of a second FastSLAM approach that relies on explicit echo detection and echo-plate boundary association. Furthermore, a SLAM experiment with a real magnetic crawler is also carried out to demonstrate the applicability of ultrasonic SLAM.

Chapter 5 presents a more elaborate mapping framework that can jointly recover a probabilistic grid representation of a metal plate, along with a feature-based representation of the same environment. The objective of this approach is to enable uncertainty assessment on the map reconstruction, and to leverage the consistency between two map representations to mitigate the detrimental effect of interference induced by DAS beamforming. Furthermore, this novel mapping framework is integrated within a frontier exploration strategy to automatically drive an autonomous agent towards informative measurement locations, and it is leveraged to extend our mapping approach for recovering more arbitrary polygonal shapes than rectangles.

In **Chapter 6**, we propose a GraphSLAM framework to solve the full SLAM problem that can leverage high-order echoes contained in the ultrasonic measurements. The proposed framework integrates a multi-order echo observation model along with DAS beamforming into a Non-linear Least Squares (NLS) optimizer. This approach is next validated on experimental measurements.

Chapter 7 presents an optimal beamforming approach for mapping a metal plate without prior knowledge of the wave propagation model, as it may be lacking during a practical inspection task. By maximizing the focusing ability of DAS beamforming, the model parameters can be adjusted with the aim to obtain more accurate mapping results when the propagation properties are not perfectly known a priori. To demonstrate the efficiency of this approach, the latter is evaluated both in nominal and non-nominal acquisition conditions.

Finally, **Chapter 8** summarizes the contributions of this thesis, and provides the reader with elements for future works.

The following table provides a summary of the different ultrasound-based localization and mapping problems that are addressed in this thesis, along with the hypotheses that are made.

Chapter/Part	Problem addressed	Online / Offline	Plate geometry	Propagation model
3	Localization	Online	Rectangular	Known
4	SLAM	Online	Rectangular	Known
5.1	Mapping	Online	Rectangular	Known
5.2	Exploration	Online	Rectangular	Known
5.3	Mapping	Offline	Polygonal	Known
6	Full SLAM	Offline	Rectangular	Known
7	Mapping	Offline	Rectangular	Unknown

Table 1.1: Summary of the problems addressed in the different chapters of this thesis.

1.4 List of publications

This thesis is built on the following articles that have been, in the majority, peer-reviewed and published either in international journals, or in the proceedings of conference venues:

- Ouabi, O. L., Pomarede, P., Geist, M., Declercq, N. F., and Pradalier, C. Monte-Carlo Localization on Metal Plates Based on Ultrasonic Guided Waves. In International Symposium on Experimental Robotics (pp. 345-353). Springer, 2020, covered in Chapter 3.
- Pradalier, C., Ouabi, O. L., Pomarede, P., and Steckel, J. On-plate localization and mapping for an inspection robot using ultrasonic guided waves: a proof of concept. In IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (pp. 5045-5050). IEEE. 2020, covered in Chapter 4.
- Ouabi, O. L., Pomarede, P., Geist, M., Declercq, N. F., and Pradalier, and C. A FastSLAM approach integrating beamforming maps for ultrasound-based robotic inspection of metal structures. IEEE Robotics and Automation Letters, 6(2), 2908-2913. 2021, covered in Chapter 4.
- Ridani, A., Ouabi, O. L., Declercq, N. F., and Pradalier, C. On-plate autonomous exploration for an inspection robot using ultrasonic guided waves. In European Conference on Mobile Robots (ECMR). IEEE. 2021, covered in Chapter 5.
- Ouabi, O. L., Pomarede, P., Zeghidour, N., Geist, M., Declercq, N., and Pradalier, C. Combined Grid and Feature-based Mapping of Metal Structures with Ultrasonic Guided Waves. In IEEE International Conference on Robotics and Automation (ICRA). (pp. 5056-5062). IEEE. 2022, covered in Chapter 5.
- Ouabi, O. L., Pomarede, P., Declercq, N. F., Zeghidour, N., Geist, M., and Pradalier, C. Learning the propagation properties of rectangular metal plates for Lamb wave-based mapping. Ultrasonics, 123, 106705. 2022, covered in Chapter 7.
- Chahine, G., Schroepfer, P., Ouabi, O. L., and Pradalier, C. A Magnetic Crawler System for Autonomous Long-Range Inspection and Maintenance on Large Structures. Sensors, 22(9), 3235. 2022.
- Ouabi, O. L., Yi, J., Zeghidour, N., Declercq, N., Geist, M., and Pradalier, C. Polygonal Shapes Reconstruction from Acoustic Echoes Using a Mobile Sensor and Beamforming. European Signal Processing Conference (EUSIPCO). 2022, covered in Chapter 5.

- Ouabi, O. L., Zeghidour, N., Declercq, N., Geist, M., and Pradalier, C. Pose-graph SLAM Using Multi-order Ultrasonic Echoes and Beamforming for Long-range Inspection Robots. *Submitted to ICRA 2023*, covered in Chapter 6.

Chapter 2

Background theory

The emergence of a complete long-range robotic inspection task based on ultrasonic waves requires the conjunction of methods from various research domains. More specifically, the adequate interpretation of the ultrasonic measurements requires that the acoustics of the wave propagation is known or, at least, modeled with sufficient accuracy. Efficient signal processing is also needed to infer, from spatially-distributed measurements, relevant information, such as the position of an acoustic reflector. Finally, conventional frameworks in robotics shall be leveraged to tackle the challenging issues, such as Simultaneous Localization and Mapping (SLAM), that arise when the target application involves a mobile robotic system.

To enable the reader to fully comprehend the guided wave-based localization and mapping techniques presented later in this thesis, this chapter introduces the fundamentals of Lamb wave theory along with the usual propagation models, the fundamentals of array signal processing using beamforming, which is a classical spatial filtering approach, and the Bayesian filtering framework, which plays a key role in robotics to solve localization, mapping and SLAM problems.

2.1 Lamb wave propagation in a metal plate

Ultrasonic guided waves (UGWs) have been identified as a promising solution to enable long-range Non-Destructive Testing (NDT) of large structure surfaces [3, 11, 86]. They can be emitted in a solid material bounded by two parallel planes using a mechanical excitation. They can propagate over large distances, with little attenuation, featuring propagation characteristics that vary with the material elastic properties. Last but not least, they are sensitive to material discontinuities, which makes their use appealing for the detection of material flaws.




Figure 2.1: Simplified representation of the out-of-plate distortion induced by Lamb wave symmetric and anti-symmetric modes in a plate. The plate boundaries at equilibrium state of the plate are represented by the discontinuous lines.




Figure 2.2: Illustration of the wave dispersion effect in a metal plate with simulated data.

Although different types of ultrasonic waves can be generated in solids, those that will be considered in this manuscript are called Lamb waves [59]. Lamb waves are peculiar in that the local deviation of the particles' position from their equilibrium state always lies in a plane that contains the propagation direction. On the contrary, Shear Horizontal (SH) waves only present particle motion that is orthogonal to the propagation direction.

2.1.1 Rayleigh-Lamb equations

The theory of Lamb waves propagation in an isotropic and homogeneous wave guide with constant thickness is well established [59]. As the particles' displacement always lies in a plane that contains the direction of wave propagation, it can be described as a combination of longitudinal and out-of-plane (transverse) motion components. Based on the wave equation and the mechanical boundary conditions induced by the wave guide, the Rayleigh-Lamb equations, which stipulate the wave modes that can exist in the solid, can be derived,

$$\frac{\tan qd/2}{\tan pd/2} = - \left[\frac{4k^2 qp}{(k^2 - q^2)^2} \right]^\alpha, \quad (2.1)$$

where d is the thickness of the plate, $k = 2\pi/\lambda$ is the wavenumber, λ is the wavelength, and the following equations define p and q :

$$p^2 = \frac{\omega^2}{c_L^2} - k^2; \quad q^2 = \frac{\omega^2}{c_T^2} - k^2.$$

In the above equation, $\omega = 2\pi f$ is the pulsation, while f is the frequency. c_L and c_T are respectively the longitudinal and transverse velocities related to the material. For $\alpha = 1$, the equation yields the physical properties of symmetric Lamb wave modes, whereas for $\alpha = -1$, the properties of anti-symmetric modes are stipulated. For symmetric modes, the displacement field is symmetric to the mid-plane of the plate, and it is anti-symmetric for anti-symmetric modes, as illustrated by Fig. 2.1. The Rayleigh-Lamb equations are solved using numerical methods [86] to determine the wavenumber k as a function of the product $f \cdot d$, or more simply, as only a function of ω , assuming a known plate thickness d . It results in a *dispersion relation* $k(\omega)$. Phase and group velocities of each mode can be subsequently determined with the relations:

$$v_P(\omega) = \frac{\omega}{k(\omega)}; \quad v_G(\omega) = \frac{d\omega}{dk} = \frac{v_P^2(\omega)}{v_P(\omega) - \omega \frac{dv_P}{d\omega}},$$

which also have a dependency in ω . This is a typical feature of dispersive propagation, which results in wavepacket distortion and spreading in the signal when the propagation distance increases. This phenomenon is illustrated in Fig. 2.2.




Figure 2.3: Dispersion curves for several low order symmetric and anti-symmetric Lamb wave modes in an aluminum plate ($c_L = 6420$ m/s, $c_T = 3040$ m/s). The top figure depicts the group velocities while the bottom plot shows the wavenumbers.

One particularity of the Rayleigh-Lamb equations is that each equation (*i.e.* Eq. (2.1) with either $\alpha = 1$ or $\alpha = -1$) always admits at least one positive real-valued solution, and the number of solutions increases with frequency. Hence, at least the two fundamental modes A_0 and S_0 (anti-symmetric and symmetric modes respectively) propagate within the material, while higher-order modes can also propagate when the excitation frequency exceeds their respective cut-off frequency. Fig. 2.3 shows the dispersion curves obtained for an aluminum material via the numerical resolution¹ of the Rayleigh-Lamb equations.

2.1.2 Green's function

The determination of a plate response to an ultrasonic excitation often requires heavy simulation through Finite Element Methods (FEM) to solve the wave equation [29]. Alternatively, one can rely on the determination (or modeling) of the Green's function g . It is defined as the spatio-temporal impulse response of a linear differential operator $L(\mathbf{x}, t)$ [40]:

$$Lg = \delta(\mathbf{x} - \mathbf{x}', t - t'),$$

so that a solution u of the inhomogeneous equation $Lu = s$, where $s(\mathbf{x}, t)$ describes the excitation applied in \mathbf{x} at time t , can be expressed as:

$$u(\mathbf{x}, t) = \int g(\mathbf{x}, \mathbf{x}', t, t') s(\mathbf{x}', t') d\mathbf{x}' dt'.$$

¹The dispersion curves are obtained using Dispersion Calculator, available at https://www.dlr.de/zlp/en/desktopdefault.aspx/tabid-14332/24874_read-61142/

A practical and computationally inexpensive approach to model Lamb wave propagation in an isotropic and homogeneous material is to rely on the Green's function related to the Helmholtz operator. Let's consider $L = \nabla^2 + k^2(\omega)$, which is the frequency-domain counterpart of the operator associated to the wave equation, and for a specific Lamb wave mode with the dispersion relation $k(\omega)$. The Green's function associated with L can be expressed using the Hankel function of 0th order and first kind $H_0^{(1)}$, yielding for an isotropic propagation around the excitation point:

$$g(r, \omega) = -\frac{j}{4} H_0^{(1)}(k(\omega)r),$$

where r is the propagation distance from the excitation. The later model is often simplified to:

$$\hat{g}(r, \omega) \approx e^{-jk(\omega)r} / \sqrt{k(\omega)r},$$

which is the Fourier transform of the impulse response. As the propagation is radial around the emitter, and the material is homogeneous, the resulting model only depends on the distance r between the emission and reception points. The displacement field is then modeled with $u(r, \omega) = \hat{g}(r, \omega) \cdot s(\omega)$. The use of this acoustic model is widespread in the literature of guided waves, especially to achieve defect detection and localization purposes [74, 43, 108, 39].

2.2 Array signal processing

Many applications involve the emission and acquisition of (electromagnetic/acoustic) signals at different locations, often using an assembly of multiple sensors, referred to as a sensor array. After data acquisition, the multichannel input measurements can be processed to infer relevant information, such as the position of an acoustic source. This section provides an introduction to array signal processing, and in particular, to beamforming, which is a traditional spatial filtering approach. Beamforming methods find various applications in Structural Health Monitoring (SHM), where ultrasonic guided waves are used by sensor networks to achieve defect detection and localization. These are briefly reviewed in the current section.

2.2.1 Introduction to beamforming

Beamforming is a versatile approach to achieve spatial filtering [97]. A signal emitted (or a reflected) by a distant object is received by multiple antenna elements arranged in an array. A beamformer can process the resulting data to achieve spatial signal selection in a way that signals coming from specific directions experience constructive interference, while signals coming from other directions experience destructive interference. It finds applications in a wide range of domains such as medical imaging [89], in wireless communications [30, 31], audio processing [41, 72], and can drastically enhance the probing capabilities of SOund Navigation and Ranging (sonar) and RAdio Detection and Ranging (radar) systems [85, 44].

For each input channel, the signal time of arrival coming from a distant source varies with the distance between the sensor and the source. A beamformer applies spatial filtering to the signals with filter weights that are characteristic of the "look" direction (or candidate source position). They are determined so that the beamformer output presents destructive interference when the look direction does not correspond to the source position, while the interference is constructive when it is the case. One of the most classical approaches is delay-and-sum beamforming. As illustrated in Fig. 2.4), the weights apply time delays to each input signal based on prior knowledge of the array geometry. Hence, source localization can be achieved by determining the location in

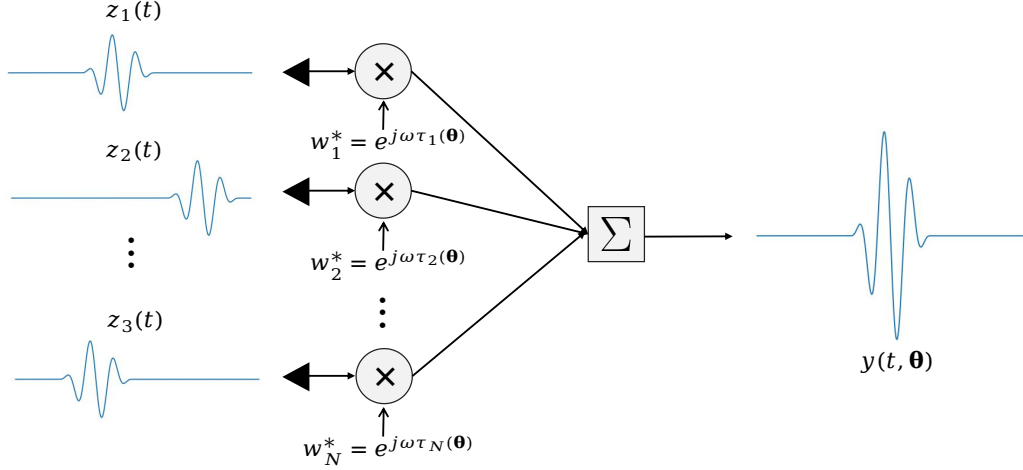


Figure 2.4: Illustration of delay-and-sum beamforming to achieve source localization. The filter weights $\mathbf{w}(\theta)$ apply delays to the input signals based on the distance between each antenna element and candidate position θ . The beamformer outputs the sum of the delayed signals which presents constructive interference when the candidate θ matches the true source location.

the search space where constructive interference on the beamformer output is maximized. Apart from the delay-and-sum approach, more advanced beamforming methods also exist, and their design is still an important research field [92].

While beamforming is often applied to fixed antennas, applications under interest increasingly involve moving sensor arrays. In synthetic aperture radar and synthetic aperture sonar for example, the sensor array is typically embedded onto a moving platform, such as an aircraft or a submarine, to probe the environment over large distances. The acquisition of data at different locations along the vehicle trajectory leads to a "virtual" antenna that can be orders of magnitudes larger than the physical one. As a result, synthetic aperture systems can achieve much higher spatial resolution than their classic version. The present thesis investigates a similar situation, except that the physical "antenna" is reduced to a single piezoelectric element: a robotic platform moving on a large metal surface continuously emits an ultrasonic excitation that generates Lamb waves in the material, and listens to the ultrasonic echoes. By processing the multichannel data with array signal processing methods, we aim at performing ultrasonic localization and mapping, with the end-objective of long-range defect detection.

2.2.2 Applications in Structural Health Monitoring

One common application of ultrasonic guided waves is Structural Health Monitoring (SHM), where piezoelectric sensors are placed in a network, and are permanently attached to the structure to be monitored (such as an aircraft nacelle [54]). To collect information about the structure integrity, ultrasonic measurements are collected by successively exciting the sensors one by one, while the others record the material acoustic response. This operation results in a data set that is to be processed to highlight the presence of a potential defect, as illustrated by Fig. 2.5.

One of the most common methods to achieve defect detection is to use baseline subtraction: the acquisition process is first performed at sensor installation (while the structure is in a healthy state), and the resulting baseline measurements are carefully recorded in memory. During the monitoring, the acquisition process is continuously performed. To each collected measurement

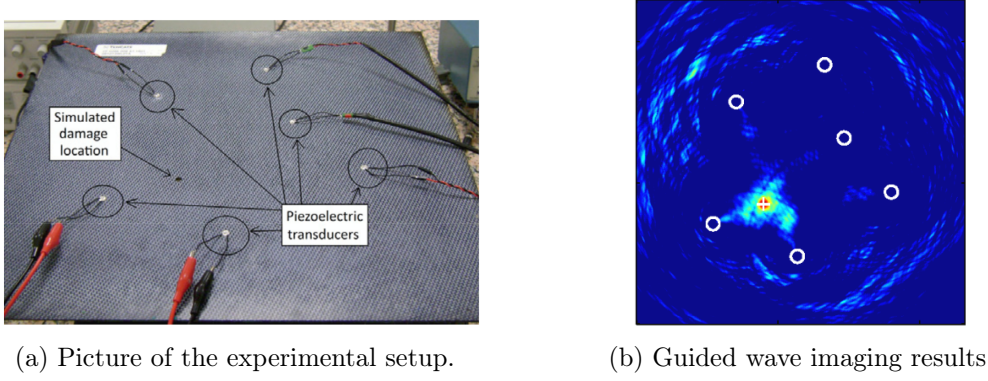


Figure 2.5: Defect detection using a network of sensors. Baseline subtraction and a multichannel processing technique are used to obtain imaging results. The illustrations are taken from [37].

related to an emitter/receiver sensor pair, the corresponding baseline measurement is subtracted. Using the resulting signals, multichannel processing methods can be applied to recover an image that highlights the presence of a potential defect [28, 37, 38]. Other methods that do not rely on baseline measurements are investigated to alleviate the introduction of errors due to varying environmental conditions [6, 60, 73, 88]. Yet, these setups can only be used to monitor a limited and fixed area, thus they are not suitable for the inspection of large metal structures.

Deploying similar methods for the long-range detection of defects on a large surface with a robotic platform is a challenging issue. In SHM, the positions of the sensors are typically measured with a high precision, which is critical for accurate defect detection results. For a robotic application however, the positions of the sensor along the robot trajectory will need to be estimated due to kinematics modeling errors, and due to the imprecise wheel displacement and rotation data provided by the wheels' encoders.

2.3 Bayesian filtering

The fact that accurate robot and acoustic reflector localization is necessary to enable long-range robotic inspection will motivate the use of probabilistic approaches to achieve meaningful inference from noisy input data. In this section, we briefly introduce the Bayesian framework, and in particular, the Bayes' recursive filter, which is conventional in robotics to achieve state estimation purposes. Next, we introduce some fundamentals on mapping, followed by an introduction on Simultaneous Localization and Mapping presented in probabilistic terms.

2.3.1 Probabilistic interaction models

State estimation through Bayesian inference requires the design of two probabilistic interaction models: a dynamic model and an observation model. The dynamic model is a distribution $p(\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{u}_t)$ that relates the robot state at time $t+1$ to the previous state \mathbf{x}_t and input data \mathbf{u}_t (odometry or control information). As we will be considering a robot moving on a planar surface, we stipulate the equations of a dynamic model relative to a robot moving in a two-dimensional world, as it will be recurrently used in the different chapters of this thesis.

The robot is assimilated to a single point. At time t , the robot state $\mathbf{x}_t = [x_t, y_t, \theta_t]^T$ contains 2D coordinates in a frame that is considered fixed to the world, along with robot heading, as illustrated by Fig. 2.6. The evolution of the robot state is governed by the dynamic model




Figure 2.6: Illustration of the robot state along with adopted notations in a 2D world.

$p(\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{u}_t)$. This formulation underlies a *Markov assumption*, because given the robot state at time t and inputs \mathbf{u}_t , we assume that all past robot states and past inputs do not bring further information on what will be the next state. As the input \mathbf{u}_t represents control commands or odometry provided by encoders, for example, they yield prior information on the translational and rotational displacement of the robot. In a noise-free situation:

$$\mathbf{u}_t = \begin{bmatrix} \Delta r_t \\ \Delta \theta_t \end{bmatrix} = \begin{bmatrix} \sqrt{(x_{t+1} - x_t)^2 + (y_{t+1} - y_t)^2} \\ \theta_{t+1} - \theta_t \end{bmatrix}.$$

Given input \mathbf{u}_t , the next state is then computed with:

$$\mathbf{x}_{t+1} = f_{\text{dyn}}(\mathbf{x}_t, \mathbf{u}_t) = \mathbf{x}_t + \begin{bmatrix} \Delta r_t \cos(\theta_t + \Delta \theta_t) \\ \Delta r_t \sin(\theta_t + \Delta \theta_t) \\ \Delta \theta_t \end{bmatrix},$$

where $f_{\text{dyn}}(\mathbf{x}_t, \mathbf{u}_t)$ is the non-linear dynamic transition model (which is often a simplification of the true robot motion). To take into account noise, we can assume that the noisy input \mathbf{u}_t is a version of the real (and unknown) robot displacement $\mathbf{u}_t^{\text{real}}$ that is disturbed by additive noise:

$$\mathbf{u}_t = \mathbf{u}_t^{\text{real}} + \begin{bmatrix} \epsilon_r \\ \epsilon_\theta \end{bmatrix},$$

where ϵ_r and ϵ_θ are samples from random zero-mean Gaussian distributions, which are respectively $\mathcal{N}(0, \sigma_r)$ and $\mathcal{N}(0, \sigma_\theta)$. Although the analytical expression of the resulting probabilistic model $p(\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{u}_t)$ can be derived, we will see that it won't be of much use in practice because sampling methods will be used to achieve Bayesian filtering.

The second interaction model necessary to perform Bayesian inference is the observation model that is used to update the state estimate based on external observations. It is represented by a probability distribution $p(\mathbf{z}|\mathbf{x})$ over some measurement \mathbf{z} given robot state \mathbf{x} . For example, given robot position and known objects, range measurements can be predicted. Yet, the perception model may be seen as a likelihood function, as the measurements are inputs, and will provide information on robot localization. For our target application, we will seek to use the reflections of ultrasonic guided waves on the boundaries of metal panels to recover accurate robot precision. The appropriate perception models to achieve this goal will be introduced in a later chapter. Overall, the dynamic process made of successive motion and perception actions performed by a robot can be represented as a directional Bayesian graph, as illustrated by Fig. 2.7, that shows the conditional dependencies between all the involved variables over time.




Figure 2.7: Representation of the dynamic process, made of successive robot motions and perception actions, with a Bayesian network. The robot poses \mathbf{x}_i are related to the past one \mathbf{x}_{i-1} and odometry input \mathbf{u}_{i-1} , while measurements \mathbf{z}_i depend only on the current state \mathbf{x}_i .

2.3.2 Recursive Bayes filter

The aim of state estimation is to determine, for any time t , the posterior distribution over the state \mathbf{x}_t based on past odometry inputs and measurements: $p(\mathbf{x}_t | \mathbf{u}_{0:t-1}, \mathbf{z}_{0:t})$, where we use the notation $\mathbf{z}_{0:t} = \{\mathbf{z}_0, \mathbf{z}_1, \dots, \mathbf{z}_t\}$. Here, we present the recursive Bayes filter, which is a widespread method to recursively determine the target posterior distribution. Based on the content from the reference [94], we recall how to derive the recursive formulation of the filter.

We seek to derive a mathematical relation between $p(\mathbf{x}_{t+1} | \mathbf{u}_{0:t}, \mathbf{z}_{0:t+1})$ and $p(\mathbf{x}_t | \mathbf{u}_{0:t-1}, \mathbf{z}_{0:t})$. We start by using the Bayes' formula:

$$\begin{aligned}
 p(\mathbf{x}_{t+1} | \mathbf{u}_{0:t}, \mathbf{z}_{0:t+1}) &= p(\mathbf{x}_{t+1} | \mathbf{u}_{0:t}, \mathbf{z}_{0:t}, \mathbf{z}_{t+1}), \\
 &= \frac{p(\mathbf{x}_{t+1}, \mathbf{z}_{t+1} | \mathbf{u}_{0:t}, \mathbf{z}_{0:t})}{p(\mathbf{z}_{t+1} | \mathbf{u}_{0:t}, \mathbf{z}_{0:t})}, \\
 &= \frac{p(\mathbf{z}_{t+1} | \mathbf{x}_{t+1}, \mathbf{u}_{0:t}, \mathbf{z}_{0:t}) p(\mathbf{x}_{t+1} | \mathbf{u}_{0:t}, \mathbf{z}_{0:t})}{p(\mathbf{z}_{t+1} | \mathbf{u}_{0:t}, \mathbf{z}_{0:t})}, \\
 &= \eta \cdot p(\mathbf{z}_{t+1} | \mathbf{x}_{t+1}, \mathbf{u}_{0:t}, \mathbf{z}_{0:t}) p(\mathbf{x}_{t+1} | \mathbf{u}_{0:t}, \mathbf{z}_{0:t}),
 \end{aligned}$$

where η is the normalizing factor of the posterior distribution (as it does not depend on \mathbf{x}_{t+1}). Based on the assumption of conditional independence, a measurement only depends on the state at the same time, while past states and inputs do not bring further information. Hence, we have:

$$p(\mathbf{z}_{t+1} | \mathbf{x}_{t+1}, \mathbf{u}_{0:t}, \mathbf{z}_{0:t}) = p(\mathbf{z}_{t+1} | \mathbf{x}_{t+1}),$$

which results in $p(\mathbf{x}_{t+1} | \mathbf{u}_{0:t}, \mathbf{z}_{0:t+1}) = \eta \cdot p(\mathbf{z}_{t+1} | \mathbf{x}_{t+1}) p(\mathbf{x}_{t+1} | \mathbf{u}_{0:t}, \mathbf{z}_{0:t})$. To explicit the recursive relation, we expand the last factor of the right part of the equation:

$$\begin{aligned}
 p(\mathbf{x}_{t+1} | \mathbf{u}_{0:t}, \mathbf{z}_{0:t}) &= \int p(\mathbf{x}_{t+1}, \mathbf{x}_t | \mathbf{u}_{0:t}, \mathbf{z}_{0:t}) d\mathbf{x}_t, \\
 &= \int p(\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{u}_{0:t}, \mathbf{z}_{0:t}) p(\mathbf{x}_t | \mathbf{u}_{0:t}, \mathbf{z}_{0:t}) d\mathbf{x}_t.
 \end{aligned}$$

We can now use the hypothesis that all the information on \mathbf{x}_{t+1} is already provided by the previous state and inputs to simplify: $p(\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{u}_{0:t}, \mathbf{z}_{0:t}) = p(\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{u}_t)$, which is the probabilistic transition model. Also, as the next input \mathbf{u}_t does not bring information on the current state

\mathbf{x}_t , we can also simplify: $p(\mathbf{x}_t|\mathbf{u}_{0:t}, \mathbf{z}_{0:t}) = p(\mathbf{x}_t|\mathbf{u}_{0:t-1}, \mathbf{z}_{0:t})$, which is our posterior at time t : $b(\mathbf{x}_t) = p(\mathbf{x}_t|\mathbf{u}_{0:t-1}, \mathbf{z}_{0:t})$. All put together, we obtain the recursive equation of the Bayes filter:

$$b(\mathbf{x}_{t+1}) = \eta \cdot p(\mathbf{z}_{t+1}|\mathbf{x}_{t+1}) \int p(\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{u}_t) b(\mathbf{x}_t) d\mathbf{x}_t. \quad (2.2)$$

Overall, the recursive Bayes filter formulae consists of two steps: a prediction step that uses the inputs \mathbf{u}_t and the belief on the state at time t to predict the next state, and an observation step that updates the estimation using an observation. Hence, a complete implementation of the filter requires three probabilistic distributions: an initial belief $b(\mathbf{x}_0)$, the stochastic transition model $p(\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{u}_t)$, and the measurement model $p(\mathbf{z}_t|\mathbf{x}_t)$. It should be noted that, when the transition and observation models are linear with respect to the state and odometry inputs, and the noise is assumed to be additive Gaussian, the posterior is also Gaussian. The mean and covariance matrix of the posterior can be iteratively determined using the well-known Kalman filter equations [94]. However, when it is not the case, approximations will be necessary, as closed-form solutions of Eq. (2.2) won't be available in practice.

It should be noted that some applications may also require the estimation of auxiliary variables to recover the state \mathbf{x} . One classic example is the estimation of correspondence variables $\mathbf{c}_{0:t}$. This situation arises when the measurements contain information on multiple map features (for example range measurements to multiple fixed landmarks), but the correspondence between each range measurement and the map feature that is being observed would not be known a priori. Adding the estimation of correspondence variables makes the state estimation problem significantly more challenging.

In our case-study, the determination of robot localization using omnidirectional ultrasonic waves may require the estimation of correspondence variables. As one measurement contains several echoes due to reflections on the plate boundaries, each echo may be associated with the boundary it originates from before robot localization can be feasible.

2.3.3 Particle filter

What follows is a mere summary of the more complete introduction on particle filters available in [94]. Particle filtering is a sequential Monte-Carlo approach that can approximate the recursive Bayes filter presented earlier. Contrary to the Kalman filter, particle filters are non-parametric. They allow for the posterior distribution $b(\mathbf{x}_t)$ to be approximated by a set of N samples (also called particles) that depend on time:

$$\mathcal{P}_t = \{\mathbf{x}_t^{(1)}, \mathbf{x}_t^{(2)}, \dots, \mathbf{x}_t^{(N)}\},$$

where each $\mathbf{x}_t^{(n)}$ represents a value in the state space. One advantage of particle filters is that they can provide an approximation of the target posterior while being relatively simple to implement (and assuming that N is large enough). Furthermore, they can be used even when the transition and observation models are non-linear, and the noise models are non-Gaussian, contrary to Kalman filters. As the Bayesian filter is a recursive approach, the particle set \mathcal{P}_{t+1} for approximating $b(\mathbf{x}_{t+1})$ is recursively determined from the previous particle set \mathcal{P}_t . The different steps for implementing the basic particle filter are as follows:

Initialization: At initialization, the set \mathcal{P}_0 is constructed by sampling N independent $\mathbf{x}_0^{(n)}$ from the designed $b(\mathbf{x}_0)$. The next steps describe the recursive update that is performed every time the input \mathbf{u}_t and observation \mathbf{z}_{t+1} are made available.

Prediction: Given the particle set \mathcal{P}_t , a new set $\mathcal{X}_{t+1} = \{\mathbf{x}_{t+1}^{(1)}, \mathbf{x}_{t+1}^{(2)}, \dots, \mathbf{x}_{t+1}^{(N)}\}$ is created by sampling, using the transition model, and for $n = 1..N$: $\mathbf{x}_{t+1}^{(n)} \sim p(\mathbf{x}_{t+1}|\mathbf{x}_t^{(n)}, \mathbf{u}_t)$.

Particle evaluation: Each particle is attributed a weight value (which is also called the importance factor): $w_{t+1}^{(n)} = p(\mathbf{z}_{t+1}|\mathbf{x}_{t+1}^{(n)})$, resulting in a higher particle weight when its state value better explains the measurement.

Particle resampling: The new particle set \mathcal{P}_{t+1} is constructed by drawing with replacement, N particles from \mathcal{X}_{t+1} , where the $\mathbf{x}_{t+1}^{(n)}$ are sampled with a probability proportional to their respective weight $w_{t+1}^{(n)}$. This results in the incorporation of measurement information in the particles' distribution.

The natural consequence of this implementation is that the particles with low weights will tend to disappear along the filtering, while particles with higher weights will tend to survive. More specifically, the particles generated by this procedure are only approximate samples of the target posterior $b(\mathbf{x}_t)$, yet, with the essential property that the empirical distribution of the samples somewhat tends to the true posterior distribution when the number N goes to infinity. Put in mathematical terms, for *any* arbitrary subset A of the state domain, the empirical fraction of samples contained in A converges to the integral of the target distribution calculated over A when the number of particles goes to infinity:

$$\frac{1}{N} \sum_{n=1}^N \mathbf{1}_{\mathbf{x}_t^{(n)} \in A} \xrightarrow{N \rightarrow \infty} \int_A b(\mathbf{x}_t) d\mathbf{x}_t.$$

The two major limitations of particle filtering lie in the fact that the naive implementation of the filter becomes greatly inefficient when the dimension of the state space becomes typically larger than 3. Another limitation is the problem of particle depletion that occurs when, due to their finite number, no particles exist in some regions of the state space, where the probability density is not null. Particle filters have arisen a lot of research interest, and various implementations have been investigated to tackle the challenges induced by their use [15].

2.3.4 Mapping

Along with robot localization, mapping is also an important problem in robotics. Mapping is as necessary as a robot needs an accurate representation of the environment to safely operate in it (for path planning and obstacle avoidance). It consists in creating a map of the environment through the determination of the posterior distribution $p(\mathbf{M}|\mathbf{x}_{0:t}, \mathbf{z}_{0:t})$ using sensor observations made at different positions. As the mapping problem only considers environmental features as unknowns to estimate, the observations positions are assumed to be flawlessly known. One reason could be that there is little uncertainty in robot kinematics and odometry measurements, or an external localization system (such as a Global Positioning System) is available.

A map is often described as a collection of objects: $\mathbf{M} = \{\mathbf{m}_k\}_{k=1..K}$. As the dimensionality K of the map can be important, a convenient approach is to assume independence between the posteriors over the individual map features. This results in the factorization [94]:

$$p(\mathbf{M}|\mathbf{x}_{0:t}, \mathbf{z}_{0:t}) = \prod_{k=1}^K p(\mathbf{m}_k|\mathbf{x}_{0:t}, \mathbf{z}_{0:t}),$$

which significantly reduces the dimensionality of the problem, as the map feature posteriors can be determined independently to each other.

In general, map representations can be regrouped in two families. A map representation can be feature-based, *i.e* it is a sparse representation made of multiple individual elements, such as anchors, that are often described by their position. Otherwise, a map representation can be grid-based: it is a dense representation that subdivides the entire space into spatial elements, each of them having a specific state, such as whether the space is being occupied by an object or not, as in the occupancy grid framework [94].

Just as the Bayesian framework provides a recursive method to achieve state estimation, it can be used to derive a recursive filter to determine the target posterior distribution over the map. A well-known example is the Bayesian binary filter with log-odd ratios [93], that can recover occupancy-grid maps from laser scans. Although this thesis will focus mainly on beamforming approaches to achieve ultrasonic mapping, we will see later, in Chapter 5, that this Bayesian filter can be used in combination with beamforming to infer complex maps from ultrasonic data.

For the same reasons presented in the case of state estimation, the estimation of correspondence variables $\mathbf{c}_{0:t}$ may also be required to solve some mapping problems. This can be particularly challenging when no prior knowledge on the number of map features that must be recovered is available.

In this thesis, we consider the problem of mapping a large metal structure using ultrasonic echoes. We will seek to map structural features such as individual plate boundaries (*i.e* weld joints), or stiffeners, that reflect the acoustic waves. Their prior identification is necessary before defect detection can be possible. Furthermore, on the robotics perspective, the detection of defect can also be considered as a mapping problem. All these elements motivate the need to elaborate advanced methods to achieve ultrasonic mapping using a mobile sensor.

2.3.5 Simultaneous Localization and Mapping

Simultaneous Localization and Mapping (SLAM) is a fundamental problem in robotics, where the objective is to construct a map of the environment concurrently to performing robot localization based on control inputs and observations. Mathematically, it consists of the determination of a joint posterior distribution $p(\mathbf{x}_{0:t}, \mathbf{M} | \mathbf{u}_{0:t-1}, \mathbf{z}_{0:t})$ over the robot trajectory and map. Due to the high-dimensionality of the {trajectory + map} space and the interdependency between the localization and mapping variables, SLAM is often a significantly more difficult problem than localization and mapping taken separately. This is all the more true when correspondence variables must be estimated as well.




Figure 2.8: Representation of online SLAM with a Bayesian network. The robot poses \mathbf{x}_i are related to the previous state \mathbf{x}_i and odometry input \mathbf{u}_{i-1} , while measurements \mathbf{z}_i depend both on the current state \mathbf{x}_i and the map \mathbf{M} .

A SLAM problem can come at two different tastes. One may seek a solution in an online fashion, which means that the posterior over the map and trajectory is recursively updated when odometry or observation inputs are made available. Fig. 2.8 shows a graphical representation of the online SLAM problem. Due to their computational efficiency, online filtering approaches are in general used for real-time robot navigation: an autonomous robot needs to continuously maintain a belief on its position and its environment to appropriately plan its next actions. Chapters 4 and 5 of this thesis will rely on this online SLAM approach. Otherwise, one may seek to solve the full SLAM problem through batch optimization using all the available measurements at once. This approach can lead to much more accurate results, but at the cost of higher computational effort. Chapter 6 of this manuscript presents a Graph-based approach to solve the full SLAM problem using ultrasonic guided waves.

For an efficient resolution of online-SLAM, an important characteristic of the problem must be highlighted. It stems from the fact that, when the robot trajectory is known, the posterior distributions over individual map features can be treated independently, as for the mapping problem. This observation motivates the following decomposition of the target posterior:

$$\begin{aligned} p(\mathbf{x}_{0:t}, \mathbf{M} | \mathbf{u}_{0:t-1}, \mathbf{z}_{0:t}) &= p(\mathbf{x}_{0:t} | \mathbf{u}_{0:t-1}, \mathbf{z}_{0:t}) p(\mathbf{M} | \mathbf{x}_{0:t}, \mathbf{z}_{0:t}), \\ &= p(\mathbf{x}_{0:t} | \mathbf{u}_{0:t-1}, \mathbf{z}_{0:t}) \prod_{k=1}^K p(\mathbf{m}_k | \mathbf{x}_{0:t}, \mathbf{z}_{0:t}). \end{aligned}$$

This formulation leads to a significant improvement in efficiency due to the low dimensionality of the map features, which can be estimated independently of each other given a robot trajectory. This approach has led to the elaboration of FastSLAM [64], which is one of the most conventional methods to solve the online-SLAM problem. FastSLAM typically uses a particle filter in the localization space, and Extended Kalman filters for estimating map features conditionally to a trajectory hypothesis [94]. In Chapter 4, we will further introduce FastSLAM, and we will propose a different FastSLAM algorithm integrating beamforming to simultaneously achieve robot localization and map the boundaries of a metal panel using ultrasonic echoes.

Chapter 3

On-plate Monte-Carlo localization

Accurate robot localization is one of the key features necessary to a true long-range robotic operation, and more generally, to any application that involves a mobile robotic platform. On a large metal structure such as a ship hull or a storage tank, ultrasonic measurements may contain the reflections of an acoustic wave (excited by the embedded ultrasonic sensor) on structural features, such as the weld joints that delimit individual metal panels. Consequently, such measurements contain position information that is to be leveraged for the target objective.

In this chapter, we propose a Monte-Carlo approach based on a particle filter to localize an emitter/receiver pair of ultrasonic sensors by relying on the reflections of ultrasonic waves on the boundaries of a metal panel. The objective is to evaluate their potential for accurate robot localization for a potential long-range inspection task on a large metal structure. The approach is tested using an isolated metal panel, in a laboratory environment. In the experiment, the acoustic measurements are acquired by moving manually the sensors on a metal surface, but the acquisition is aimed to be performed by a robot (which will be done later in this thesis). The results demonstrate that, on the considered plate sample, sensor localization is possible with an accuracy of a few millimeters.

3.1 Introduction

The accurate localization of a mobile platform is critical for a successful long-range inspection task. Indeed, such an application may require the operational robot, such as a magnetic crawler, to reliably navigate on the structure, so that it can accomplish its mission such as reaching a target measurement location while avoiding obstacles. This requires that robot position is accurately known at any time of the operation. Besides, accurate robot positioning is required for the purpose of long-range defect detection: as measurement positions are inputs to acoustic imaging techniques [86], their precise knowledge is critical for accurate inspection results.

Localization on a metal structure may be achieved by using Lamb waves: a mobile robot equipped with a piezoelectric acoustic sensor would continuously excite an omnidirectional ultrasonic wave, which would radially propagate in the metal structure over large distances. By listening to the ultrasonic echoes due to reflections on structural features, such as plate boundaries, one may achieve robot localization by relying on the acoustic reflectors as structure-bound landmarks. Furthermore, in combination with other localization systems (Inertial Measurement Unit (IMU), Ultra-wide band beacons, laser theodolite, for example), ultrasonic measurements could be used for accurate positioning. In this chapter, we address the localization problem for a robot moving on a rectangular metal panel, which represents the constitutive element of large

metal structures such as ship hulls and storage tanks. We elaborate an approach based on a particle filter to enable the accurate localization of an industrial robot on a large metal structure using the ultrasonic reflections of the excited wave on the plate boundaries. Robot localization from external measurements based on a particle filter is typically referred to as Monte-Carlo localization [22].

3.1.1 Related works

Robot localization is a problem that is well studied in robotics. Among the standard solutions, a class of localization systems are based on preinstalled infrastructures, and typically rely on direct-path range measurements to some known and fixed landmarks for positioning. GPS, ultra-wide band beacons [36] and sound beacons for underwater localization [107] provide instances of such systems, and are deployed for navigation in outdoor and clutter-free environments.

However, when considering more challenging situations, such as navigation in an enclosed space, one may seek a localization system which does not rely on preinstalled infrastructures, as they can be impractical to use, or are simply unsuitable. In such situations, robot localization may be achieved by relying directly on the sensor suit embedded on the mobile platform. For example, distances to obstacles, such as the walls of a room, can be obtained from measurement scans provided by an embedded lidar or sonar, and the resulting measurements can be integrated within a probabilistic localization algorithm [22]. Localization in an outdoor environment on a large metal structure is expected to be a challenging issue, as the typical embedded sensors (camera, laser range finder, sonar) may be impractical to use due to the lack of consistent structural features for robust localization.

It is envisioned that, within a data fusion framework, omnidirectional UGWs can be used, in combination with appropriate sensors (IMU especially), for accurate localization on a large metal surface by relying on the ultrasonic reflections on structural features such as plate boundaries. In the literature, UGWs are often deployed in static networks of sensors for long-range defect detection and localization in a fixed inspection surface [37, 73, 86, 67]. In this setup, the sensors are permanently attached to the structure, and their positions are known precisely. Conversely, data acquisition with a sensor embedded on a mobile platform leads to flaw in the localization of the measurement points due to noisy odometry information, noisy control inputs or kinematic modeling errors. If not corrected, positioning errors may lead to poor inspection results through acoustic imaging techniques [23].

Localization based on acoustic reflections has been the subject of various studies in the literature. In [84], a rectangular environment is assumed, with known sizes, and the focus is on localization. However, the idea of guided wave-based localization has been little explored in the literature. A localization approach relying on ultrasonic reflections on a metal plate boundaries is proposed in [99], but it makes important assumptions, such as relying on known correspondence between the echo measurements and the edges they originate from. Besides, the work is restricted to numerical validation. Hence, the practical localization of a robot on a metal structure using Lamb waves is yet to be thoroughly investigated.

In the following, we elaborate a Monte-Carlo approach to evaluate the potential of ultrasonic wave-based on-plate localization. The most important challenges of such an application arise due to the omnidirectionality of the wave and the multiplicity of the ultrasonic reflections. Hence, not only the individual echoes within the measurements need to be detected, but they also need to be associated with the reflector they originate from. Recent works in the literature address similar issues for acoustic localization [78, 8, 100], where acoustic echoes need both to be detected and labelled. Contrasting with such approaches that are often complex to implement, and yield

localization results sensitive to detection errors, we propose a simple Monte-Carlo approach that does require explicit echo detection and association.

We will show that by leveraging the waves reflections on the boundaries of a rectangular panel and a basic wave propagation model, a particle filter can estimate the position of a moving emitter/receiver ultrasonic sensor despite the complex nature of the guided waves (which are multi-modal and present dispersive propagation), the sensing symmetries intrinsic to the plate's geometry and the highly reverberating propagation media. The performance of our approach is assessed with real acoustic measurements that are acquired by hand on a relatively small metal plate in a laboratory environment. The experimental results demonstrate that in this scenario, the sensor position can be recovered with a millimeter precision.

3.2 On-plate localization using ultrasonic guided waves

We propose a method based on a particle filter for the localization of an industrial robot on a large metal structure that leverages first-order ultrasonic reflections of an excited wave on a metal plate's edges. The overall objective is to evaluate the potential of UGWs for accurate sensor localization. We specifically use a wave propagation model that is correlated to the measurement with the objective to design the observation model of the filter. The specificity of such an approach is that echo detection and echo association, *i.e* the determination of which echo corresponds to a reflection on which plate boundary, are not explicitly solved. This results in a simple and easy to implement algorithm that may be more robust to noise. Besides, as only the localization problem is addressed here, the proposed method assumes that the dimensions of the metal plate are known.

3.2.1 Problem statement

We are considering a mobile unit equipped with a co-located emitter/receiver pair of transducers, and moving on a metal panel that we assume to have a rectangular shape with known dimensions $w \times h$. This constraint on the plate geometry may not be seen as a limitation for localization on storage tanks or ship hulls, as they are almost entirely made of rectangular panels. At the i^{th} scanning position, the emitter sends a pulse $s(t)$ to excite guided waves in the plate material, and the receiver collects the acoustic response $z_i(t)$, which contains the ultrasonic echoes on the plate edges. This aims at approximating sonar-like pulse-echo sensing, though constrained to an acoustic wave guide. Our objective is to recover, for any time step t , an estimate $p(\mathbf{x}_t | \mathbf{u}_{0:t-1}, z_{0:t})$ of the robot state $\mathbf{x}_t = [x_t, y_t, \theta_t]^T$ (2D Cartesian coordinates and heading) from the acoustic measurements and noisy odometry inputs $\mathbf{u}_t = [\Delta r_t, \Delta \theta_t]^T$ (translation and rotational movements). The estimation is done in the metal panel frame, *i.e.* the origin O is situated in the bottom left corner of the plate. The x and y axis are respectively aligned with the edges of size w and h . An illustration of the adopted notations is provided in Fig. 3.1.

3.2.2 Measurement model and image sources

In our setup, the acoustic measurements essentially consist in a succession of the reflections of the excited wave on the plate boundaries. As the small-sized corrosion patches one might aim to detect with robotic inspection may only act as tiny acoustic scatterers, their potential effect can be neglected in first approximation. Under the assumption that the physics of wave propagation are governed by linear dynamics, the received measurements can be modeled with:

$$z_i(t) = g_{\mathbf{x}_i}(t) * s(t) + n_i(t),$$




Figure 3.1: Illustration representing a robotic agent whose pose is $\mathbf{x} = [x, y, \theta]^T$ on a rectangular metal plate of dimensions $w \times h$.

where $g_{\mathbf{x}_i}(t)$ is the plate acoustic response received in \mathbf{x}_i for an impulse signal being emitted at the same position, and $*$ denotes the convolution operation. n_i is assumed to be an additive Gaussian noise term. To further explicit the structure of the signal model, we can assume that the material is perfectly isotropic, the basic laws of signal reflection hold to describe echoes from the plate edges, and mode conversion does not occur due to wave reflection. Under these assumptions, a standard measurement model to reverberation is the image source model [58]. It relies on the fact that each received echo signal due to a reflection from a plate boundary can be considered as a signal originating from a fictional source. The position of the image sources can be deduced from the real source position and the reverberant media geometry. When considering a reverberant media with a polygonal shape, first-order image sources can be constructed by symmetrically reflecting the real source on the media boundaries. High-order image sources (which are due to multiple reflections on different boundaries) can be recursively determined by reflecting the image sources with a lower order on the boundaries [58], as illustrated by Fig. 3.2. The use of the image source model results in the following measurement model:

$$z_i(t) = \sum_{\mathbf{p} \in \mathcal{I}(\mathbf{x}_i)} \lambda_{i,\mathbf{p}} \cdot g(\|\mathbf{p} - \mathbf{x}_i\|, t) * s(t) + n_i(t),$$

which decomposes the measurement into individual signals emitted by image sources. In the above equation, $\mathcal{I}(\mathbf{x}_i)$ is the set of image sources positions when the real source is in \mathbf{x}_i , the notation $\|\cdot\|$ refers to the Euclidean distance calculated over the Cartesian x, y coordinates, and $g(\|\mathbf{p} - \mathbf{x}_i\|, t)$ is the acoustic response of the infinite plate to an impulse being generated in \mathbf{p} and received in \mathbf{x}_i . Due to the isotropic assumption, the propagation model only depends on the travel distance between the fictive image source and the receiver. $\lambda_{i,\mathbf{p}}$ is the amplitude of the signal originating from the image source \mathbf{p} . The exact $\mathcal{I}(\mathbf{x}_i)$ can be determined for arbitrary polygonal shapes using the image source algorithm [9], but it won't be used here.

For simplicity, our method will only rely on first-order echoes, *i.e.* wave reflections that are received after exactly one reflection on a plate edge. We thus decompose the image source set as

$$\mathcal{I}(\mathbf{x}_i) = \mathcal{I}_{\text{1st-order}}(\mathbf{x}_i) \cup \mathcal{I}_{\text{high-order}}(\mathbf{x}_i).$$




Figure 3.2: Illustration of the image source model. On the left is an actual setup with an omnidirectional emitter/receiver situated in a space partially bounded by two reflective lines. On the right is the equivalent setup provided by the image source model. The first-order image sources at \mathbf{p}_1 and \mathbf{p}_2 are obtained by symmetry of the real source on each boundary. The second-order one at \mathbf{p}_3 is obtained by symmetry of one of the first-order image sources.

This results in the measurement model:

$$\begin{aligned}
 z_i(t) &= \sum_{\mathbf{p} \in \mathcal{I}_{1\text{st-order}}(\mathbf{x}_i)} \lambda_{i,\mathbf{p}} \cdot g(\|\mathbf{p} - \mathbf{x}_i\|, t) * s(t) \\
 &\quad + \sum_{\mathbf{p} \in \mathcal{I}_{\text{high-order}}(\mathbf{x}_i)} \lambda_{i,\mathbf{p}} \cdot g(\|\mathbf{p} - \mathbf{x}_i\|, t) * s(t) + n_i(t), \\
 &= \sum_{j=1}^4 \lambda_{i,j} \cdot g(d_j, t) * s(t) + \tilde{n}_i(t),
 \end{aligned}$$

where \tilde{n}_i regroups high-order echo signals and the initial noise term n_i . This notation highlights that wave reflections that are received after more than one reflection are considered as noise. The values d_j represent the distances to first-order image sources, which correspond to twice the orthogonal distance to each plate boundary. For a rectangular plate with dimensions $w \times h$ and an emitter/receiver at $\mathbf{x} = [x, y, \theta]$, these distances can be easily determined: $[d_1, d_2, d_3, d_4] = [2x, 2y, 2(w - x), 2(h - y)]$.

3.2.3 Correlation signal

Echo detection from ultrasonic measurements is necessary prior to achieving robot localization, as they contain crucial information on the ranges to the plate boundaries. With the aim to retrieve these distances from a measurement $z_i(t)$, we rely on a wave propagation model \hat{g} to estimate the likelihood that a reflection of the excited wave occurred at a distance r from the sensors. First, we consider the signal that would only contain such a reflection: $\hat{z}(r, t) = \hat{g}(2r, t) * s(t)$, where \hat{g} is the impulse response of the supposedly isotropic material. When considering Lamb waves, we will always make the assumption that the excitation signal is properly chosen, so that there is a single predominant A0 mode that propagates (while the others can be neglected, or considered as noise). Next, we rely on the acoustic model presented in the previous chapter:

$$\hat{g}(r, \omega) \approx e^{-jk(\omega)r} / \sqrt{k(\omega)r}.$$

where $k(\omega)$ is the wavenumber of the major acoustic mode that can be determined by solving the Rayleigh-Lamb equations given prior knowledge of the mechanical properties of the plate material. We then create a new signal through correlation of the measurement with the signal expected after reflection on a boundary at a distance r :

$$z'_i(r) = \frac{\langle z_i(t), \hat{z}(r, t) \rangle}{\sqrt{\langle z_i(t), z_i(t) \rangle \langle \hat{z}(r, t), \hat{z}(r, t) \rangle}}, \quad (3.1)$$

where $\langle \cdot, \cdot \rangle$ denotes the scalar product in the domain of continuous signals: $\langle u(t), v(t) \rangle = \int_{-\infty}^{+\infty} u(\tau)v(\tau)d\tau$. The correlation is expected to be high when a reflector is indeed present at a distance r . As the resulting signal z'_i presents oscillations consistent with the wave spatial periodicity, it is more convenient to only work with its envelope, which may not present large drops in value for slight variations in r :

$$e_i(r) = |z'_i(r) + j\mathcal{H}(z'_i)(r)| \quad (3.2)$$

where \mathcal{H} denotes the Hilbert transform operator. Hence, the resulting signal e_i takes its values only between 0 and 1, and a higher value at r translates into a high likelihood that a reflection occurred at such a distance. Overall, by looking at the local maxima of the envelope signal $e_i(r)$, one can deduce the most likely reflections. It should be noted that high-order reflections may also yield local maxima. This processing method will be recurrently used in the next chapters of this thesis to solve a variety of ultrasonic localization and mapping problems of interest.

3.2.4 Monte-Carlo localization

We propose a particle filter approach to recover the position of a mobile sensor on a metal plate. As we will rely on a basic implementation of the filter, the current objective is to design an observation model $p(\mathbf{z}|\mathbf{x})$ that can leverage the multiple ultrasonic echoes contained in the measurements in a suitable manner.

We propose to design a simple observation model based on the envelope signals e_i , which does not require explicit echo extraction and association. The main idea is that the measurement likelihood should be high when the likelihood of observing each of the four expected first-order reflections is also high. Based on an independence assumption between the observation of the individual echoes, a simple model that can achieve the target property is:

$$p(\mathbf{z}_i|\mathbf{x}) = \eta \cdot \exp \left\{ \beta (e_i(x) + e_i(y) + e_i(w - x) + e_i(h - y)) \right\}, \quad (3.3)$$

where β is a strictly positive parameter. This model will be used to calculate the importance weights of the particle filter, so that particles that hold a position hypothesis that is more consistent with the observation of first-order reflections will tend to survive, while the others will tend to disappear. For the localization problem, each particle represents a belief on the robot state at any measurement step i : $\mathbf{x}_i^{(n)} = [x_i^{(n)}, y_i^{(n)}, \theta_i^{(n)}]^T, n = 1, 2, \dots, N$. Our implementation of the filter using ultrasonic echoes is then the following:

Initialization: The particle filter is initialized by sampling N particles on the metal plate using uniform distributions. The next steps are performed every time some displacement information \mathbf{u}_i and measurement \mathbf{z}_{i+1} are made available.

Prediction: The motion model is applied to every particle:

$$\mathbf{x}_{i+1}^{(n)} \sim p(\mathbf{x}_{i+1}|\mathbf{x}_i^{(n)}, \mathbf{u}_i).$$

To achieve this in practice, translation and rotation displacements are first sampled with $\Delta r_t^{(n)} \sim \mathcal{N}(\Delta r_t, \sigma_r)$ and $\Delta \theta_t^{(n)} \sim \mathcal{N}(\Delta \theta_t, \sigma_\theta)$. Next, the new particle state is inferred with:

$$\mathbf{x}_{t+1}^{(n)} = f_{\text{dyn}}(\mathbf{x}_t^{(n)}, \mathbf{u}_t^{(n)}) = \mathbf{x}_t^{(n)} + \begin{bmatrix} \Delta r_t^{(n)} \cos(\theta_t^{(n)} + \Delta \theta_t^{(n)}) \\ \Delta r_t^{(n)} \sin(\theta_t^{(n)} + \Delta \theta_t^{(n)}) \\ \Delta \theta_t^{(n)} \end{bmatrix},$$

If one limits the prediction to this single step, the particle filter might be likely to fail when the particles have converged in a restricted area of the state space that is not in the vicinity of the true state, as there would be no mechanism to recover from such a situation due to the sampling scheme. This issue is well known under the name of particle depletion [94]. A simple remedy to this problem consists in disturbing the so-called proposal distribution [94] to generate particles further away in the state space. Here, we randomly disturb a particle state by drawing, after the prediction and with a low probability γ , a new sample:

$$\mathbf{x}_{i+1}^{(n)} \leftarrow \mathcal{N}(\mathbf{x}_{i+1}^{(n)}, \Sigma),$$

where \mathcal{N} represents the normal distribution and Σ is a parameter of the filter.

Update and resampling: We compute, for every particle, the importance weight by using the observation model defined in Eq. (3.3):

$$w_{i+1}^{(n)} = p(\mathbf{z}_{i+1} | \mathbf{x}_{i+1}^{(n)}).$$

Finally, we create a new set of particles by drawing the n -th particle with a probability proportional to $w_{i+1}^{(n)}$. This procedure is then repeated for the next measurement step. Due to the low computational complexity of the approach, it can be seamlessly run in an online fashion. One may also note that no explicit echo detection and association is required, which results in a simple algorithm implementation.


It should be noted that, due to the disturbance of the proposal distribution, the convergence property of the particle filter no longer holds. More specifically, the particle's distribution does not converge, when N goes to infinity, to the target posterior distribution provided by the Bayes' filter formula if $\gamma > 0$. However, one may expect that it is sufficiently close to it provided that γ is sufficiently small. Besides, the loss of this theoretical property is widely compensated by the improved localization results that can be obtained due to the ability of the filter to recover from a poorly distributed particle set. This will be illustrated in the results section.

3.3 Experimental results


In this section, we assess the performance of our approach using experimental ultrasonic measurements acquired on a metal panel, in a laboratory environment. The data are acquired by manually placing the transducers, but the acquisition is aimed to be performed by a robotic platform. The experimental results demonstrate that, in the considered scenario, the sensor position can be recovered with a precision of a few millimeters by relying on the ultrasonic echoes.

3.3.1 Experimental setup

We acquired acoustic measurements by placing manually nearly collocated contact piezoelectric V103-RM U8403008 transducers, used as an emitter/receiver pair, at predetermined positions



(a) Experimental setup



(b) Acquisition positions

Figure 3.3: (Left) Picture of the experimental setup for data acquisition. (Right) Position of the measurement points on the plate.

on a metal plate. It is to be noted that these transducers are designed to generate 1MHz signals, which is a much higher frequency than these we will commonly use. Yet, these sensors are used here because, at the time of the experiment, it was the best solution available to generate and receive Lamb waves. Although it may not be an optimal choice, the resulting signal-to-noise ratio was reasonably satisfying. The metal plate used for the experiment has dimensions $w \times h = 60 \times 45$ cm, is made of aluminum, and has a thickness of 6 mm. It also contains small holes of different sizes and depth. A picture of the experimental setup can be seen in Fig. 3.3a.

To generate Lamb waves in the plate material, we use an excitation signal that corresponds to two tone-bursts of a sinusoidal waveform at 100 kHz, generated with a USB-6356 National Instruments data acquisition (DAQ) device. The signal is then amplified to up to 100 V peak-to-peak with an analog amplifier before being sent to the emitter, while the receiver collects the ultrasonic response which contains the acoustic echoes. During each acquisition, we use coupling gel to ensure good contact between the transducers and the plate surface. Also, we numerically remove the incident wave packet from the measurements with a smooth sigmoid window. This signal is due to the direct transmission of the excitation to the receiver, thus it does not correspond to a reflection from a plate boundary. In total, we record 108 measurements for as many acquisition positions, on a regular grid, as depicted in Fig. 3.3b. Each measurement contains $M = 500$ samples collected at a sampling rate of 1.25 MHz.

One may note that, in our experiment, two separate transducers are used, whereas our method is based on a point-like and colocated emitter/receiver assumption. Here, the latter is assimilated to the center between the two transducers. This may not cause a significant error, as the diameter of the transducers (1.7cm), and thus the distance between their centers, is not large compared to the wavelength (2 cm) for the considered excitation frequency. Thus, the difference induced by this setup has little impact on the signals.

3.3.2 Dispersion curves

An important feature of our localization approach is that it is *model-based*. More specifically, it explicitly requires a model of the wave propagation, and the adequacy of the latter with the measurements will be central for accurate results.

To validate the assumptions made on wave propagation, and in particular, the assumption




Figure 3.4: Two-dimensional space-time Fourier transform of the experimental wave field recorded on the considered aluminum plate sample. It is overlaid by the theoretical dispersion curves of the A0 and S0 Lamb wave modes.

that the A0 mode is predominant, we experimentally evaluate the matching between the designed propagation model and the experimental signals. For this purpose, we experimentally determine the dispersion curves of the wave field through numerical calculation of the 2 dimensional space-time Fourier transform of the incident signal, which is theoretically given by:

$$\tilde{z}(\omega, k) = \int \int z(t, r) e^{-j(\omega t + k r)} dt dr.$$

In the above equation, the wave field $z(t, r)$ is obtained by recording the experimental time response of the incident wave for several distances r from the emitter regularly distributed, on the aluminum plate. In Fig. 3.4, we show the experimental amplitude $|z(\omega, k)|$ obtained using an excitation signal centered at 100 kHz, and a spatial periodicity of 2.2 mm between each measurement. We also display the theoretical dispersion curves of the A0 and S0 modes obtained by solving the Rayleigh-Lamb equations. It can be observed that, in our setup, the propagation model fits the measurements as there is a good match with the theoretical A0 curve. Also, the A0 mode is predominant, as little energy follows the S0 theoretical curve. Overall, the use of an A0 mode-only propagation model appears legitimate to achieve echo detection.

3.3.3 Echo detection

In this subsection, we illustrate the echo-detection principle based on the correlation signals. An acoustic measurement is acquired by placing the emitter/receiver transducer pair at 8 cm to an edge, in a corner of the plate, as depicted by Fig. 3.5. In Fig. 3.6, we show the resulting measurement (top), and the correlation signal obtained using Eq. (3.1), along with its envelope calculated with Eq. (3.2). (bottom). The envelope is the signal which is fed to the particle filter. It can be seen in Fig. 3.6 bottom, that we manage to retrieve, from the position of the local maxima, all the distances where first-order reflections occurred which are 8, 37 and 52 cm. The echo detected at nearly 45 cm corresponds to a higher-order reflection, and it has an




Figure 3.5: Illustration showing the position of the sensors on the metal plate for acquisition of the test measurement.




Figure 3.6: The acoustic measurement acquired on the aluminum plate after placing the transducers on the plate, as depicted in Fig. 3.5 (top plot), and the correlation signal obtained from the measurement along with its envelope (bottom plot). The ranges relative to the first-order reflections (8, 37, 52 cm) can be successfully retrieved from the local maxima of the envelope signal. The echo at nearly 45 cm corresponds to a higher-order reflection.

amplitude that is comparable to that of the first-order wave packets (if not higher). Such high-order reflections are considered as noise by our approach. Hence, we will determine a posteriori if their presence has a detrimental effect on the results.

3.3.4 Localization results

To test our framework, we simulate a robot following a "lawn-mower" trajectory along the y-axis by presenting the corresponding sequence of measurements to the particle filter, along with the theoretic displacement of the transducers between two measurements as odometry data. To simulate noisy inputs, we add Gaussian noise on the ground true odometry information: $\tilde{\Delta}r \sim \mathcal{N}(\Delta r, (10^{-2}\Delta r + \Delta r_0)^2)$ and $\tilde{\Delta}\theta \sim \mathcal{N}(\Delta\theta, (10^{-2}\Delta\theta + \Delta\theta_0)^2)$ with $\Delta r_0 = 10^{-3}$ m and

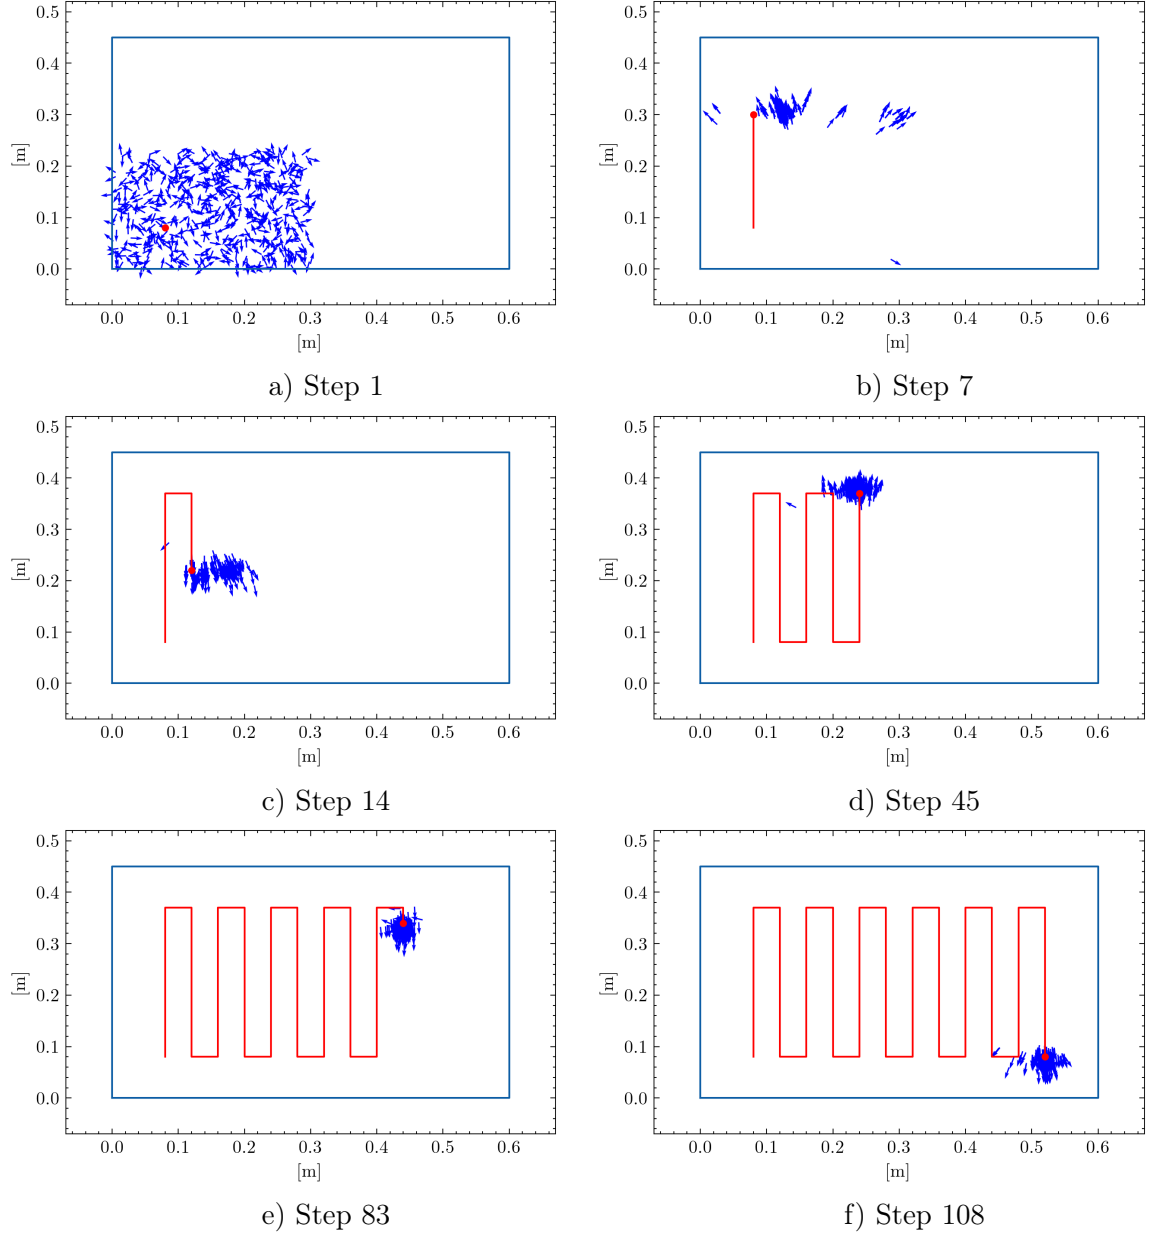


Figure 3.7: Qualitative localization results of the particle filter during measurement steps 1, 7, 14, 45, 83 and 108 of the simulated lawn mowing trajectory. The red dot represents the ground truth sensor position, the red line accounts for the true sensor trajectory. The blue rectangle represents the outline of the plate, and the arrows account for the particles' distribution.




Figure 3.8: Evolution of the localization errors (determined based on the median of the particles' state) on the x (left) and y (right) coordinates for 10 repetitions of the considered scenario. The red line is the average error calculated over the repetitions.




Figure 3.9: Evolution of the localization errors (determined based on the median of the particles' state) on the x (left) and y (right) coordinates for 10 repetitions of the considered scenario, and using the particle disturbance probability $\gamma = 0$. The red line is the average error calculated over the repetitions.

$\Delta\theta_0 = 10^{-2}$ rad. With such noise level, we aim at simulating moderate odometry uncertainty due to the robot magnetic adherence and embedded accelerometers used to provide precise heading on a nearly-vertical structure, in a real-world scenario.

The particle filter is then run for the designed scenario using $N = 500$ particles. The filter parameters used in our trials are $\beta = 5$, $\gamma = 0.03$ and $\Sigma = \text{diag}(0.01, 0.01, \pi/10)$. We first illustrate the performance of the designed approach in Fig. 3.7, which depicts the position and orientation of the particles during the measurement steps 1, 7, 14, 45, 83 and 108 of the simulated lawn-mower trajectory. To ensure the uniqueness of the position estimate despite the plate symmetries, the particles have been initialized randomly in the bottom-left part of the plate. During the first step, as both the position and orientation of the fictive robot are very uncertain, the particles have been sampled from uniform distributions. During measurement steps 7 and 14, the particles' distribution becomes more concentrated in the vicinity of the true sensor position. At Step 45, the particles have converged in the close area around the ground true position, and it can be observed that this is still the case until the very last measurement step, despite the disturbance of the proposal distribution.

To quantitatively assess the accuracy and repeatability of the estimation, we show, in Fig. 3.8, the evolution of the estimation errors on the x and y coordinates along the trajectory, for a 10-run Monte-Carlo simulation of the scenario under consideration. To consider point estimates that are little sensitive to out-of-distribution particles (as some are generated on purpose to alleviate particle depletion), the errors are defined as the absolute value of the difference between the median coordinate calculated over the particles' state, and the ground truth coordinate. The results show that, after a limited number of steps, an estimation accuracy better than 1 cm is seamlessly achieved both on the estimation of x and y , and it is the case for all the repetitions.

To further highlight the benefit of our approach which relies on the disturbance of the proposal distribution, we carry out the same experience as previously without this specific step (*i.e.* using a value $\gamma = 0$). The estimation results depicted in Fig. 3.9 show that, although the estimation accuracy appears similar for the estimation of y compared to the previous experiment, the estimation errors on x are seemingly higher. There is especially one repetition for which correct estimation convergence has not been fulfilled by the end of the trajectory. This result highlights the benefit of proposal disturbance to improve the filter convergence.

Overall, the results show that the use of UGWs combined with our particle filter approach has potential for accurate estimation of 2D coordinates on a metal surface. One may note that, in our experiments, it may not be possible to further improve the localization precision due to the model's inaccuracy, as the latter only considers a single collocated emitter/receiver.

3.4 Chapter summary

This chapter aims at demonstrating the potential of omnidirectional ultrasonic guided waves for accurate robot localization on a metal structure by relying on the acoustic reflections on a metal plate boundaries. In the considered setup, the acoustic emitter/receiver is moving on the metal surface, and only noisy information on the sensor displacement and the dimensions of the plate are assumed to be available for localization.

To achieve the target objective, we have proposed a simple Monte-Carlo approach based on a particle filter. First, correlation signals are determined by correlating the measurements with a model of the wave propagation, so that the distances between the sensor and the plate boundaries can be retrieved. These signals are subsequently integrated into an observation model that leverages the reception of first-order echoes. Next, a basic implementation of the particle filter is proposed. It features a two-step resampling process in the prediction step to maintain a high particle diversity, with the aim to alleviate particle depletion. The overall method is relatively simple to implement and to set up: along with the transition model that can be tuned based on a prior belief on the odometry noise level, only the parameters β , γ and Σ need to be specified to run the particle filter. An important particularity of this approach is that it does not necessitate explicit echo detection or echo association with the plate boundary it originates from, which often results in a complex algorithm sensitive to noise.

We have tested our approach on a $600 \times 450 \times 6$ mm aluminum metal plate sample, in a laboratory environment. To acquire the acoustic data, a nearly collocated emitter/receiver pair of contact piezoelectric transducers is moved manually to predetermined positions. The trajectory of a robot is subsequently simulated using the theoretic displacement between the measurement positions disturbed with additive noise as odometry.

The results demonstrate the potential of ultrasonic guided waves with our proposed particle filter approach. More specifically, millimeter accuracy on localization is achieved in the considered experimental setup. Despite the limitations inherent to the experimental conditions that are not representative of a realistic scenario, it is envisioned that, on a large metal structure, ultrasonic reflections could be leveraged by a mobile robot for accurate positioning, which is critical to long-range defect detection through acoustic imaging techniques.

Chapter 4

Online Simultaneous Localization and Mapping

This chapter focuses on Simultaneous Localization and Mapping (SLAM), which is a more significant problem than localization only. By relying on the reflections of UGWs on the boundaries of a metal plate, we seek to recover both the sensor trajectory and the location of the plate edges. The chore of the proposed solution is FastSLAM, which is a particle filter-based approach that relies on a factorization of the joint target posterior over the trajectory and map to efficiently solve the online SLAM problem. Its essence is to use a particle filter, where each particle holds a hypothesis on the full trajectory and a map inferred from the trajectory.

In the following, two FastSLAM approaches are reviewed. The first method has been elaborated before the start of this PhD thesis. It relies on a simple signal model, and attempts to achieve explicit echo detection and association with estimated line features, within the FastSLAM framework. The second approach is a contribution of this thesis: it leverages Lamb wave theory to take into account wave dispersion in the measurement model, and beamforming to localize plate boundaries. The merits of this approach are that explicit echo detection and association, which are challenging issues when considering ultrasonic signals, are no longer required. These elements are integrated within a FastSLAM algorithm.

The two approaches are subsequently compared using experimental data acquired on isolated metal panels, in a laboratory environment. The results show that the novel approach provides more accurate localization and mapping results with less algorithmic complexity. Yet, it leverages prior knowledge of the rectangular plate geometry, contrary to the former approach. Finally, a demonstration using a magnetic crawler is carried out to illustrate the feasibility of deploying the proposed ultrasonic SLAM framework with a real robotic platform.

4.1 Introduction

As presented in the previous chapter, precise localization of the robot on a large metal structure is necessary before long-range defect detection can become feasible. Robot localization typically relies on external measurements that can be used to relate robot position to structural (and usually fixed) landmarks. In the previous chapter, it has been demonstrated that the ultrasonic reflections on the edges of a metal plate can be leveraged for this purpose. Yet, the proposed localization framework relies on prior knowledge of the plate geometry.

The possibility to build a map of the environment (*i.e* to estimate the position of plate boundaries in our case-study) is a desired objective. In realistic scenarios, for example, a large

metal structure may be made out of plates with different sizes, and a precise map depicting the position of the weld joints on the whole structure may not be readily available. Another reason to aim for that objective is that detection of defects (which may also act as acoustic scatterers) shall be considered as a mapping problem as well, and their localization is the end-objective, along with robot positioning. Overall, simultaneous localization and mapping is a fundamental problem that must be solved to enable the emergence of long-range robotic inspection.

Apart from the application of interest, acoustic SLAM methods can find a broad range of applications in different environments. Examples of such may include marina underwater environments such as dams or harbors, and using a sonar [76, 17], enclosed spaces using a set of distributed microphones [24, 50], or using a single omnidirectional emitter/receiver [51, 53]. It is envisioned that the methods proposed to solve the considered SLAM problem could be beneficial for such related applications.

Before diving into ultrasonic SLAM, we will briefly review the literature on acoustic mapping and the challenges related to the use of omnidirectional UGWs.

4.1.1 Acoustic mapping

The mapping problem is fundamental in robotics, as a robot typically needs a map of the environment to safely navigate in it. Also, when integrated within a SLAM framework, the recovered map can be used for accurate robot localization. Some of the most traditional methods are dedicated to mapping indoor environments such as offices for example, and using sensors that emit a signal and receive its reflection on an object in a known direction, such as LIDARs and sonars [26, 65]. Due to the directivity hypothesis on the sensor models, traditional mapping methods are unsuitable when relying on omnidirectional ultrasonic waves.

Apart from defect detection purposes, UGWs have been used to map structural features such as the edges of a metal plate [82], or stiffeners [42], by relying on scattered signals and appropriate signal processing. Yet the mentioned works consider static sensor networks. On the other way, several studies [90, 91] have considered embedding an ultrasonic sensor onto a moving robotic platform for mapping a large structure, but they rely on directional guided waves. Hence, only little work has considered the joint use of omnidirectional UGWs and moving sensors. [62] proposes a method for mapping plate boundaries, while [32] evaluates the feasibility of mapping other structural features such as stiffeners. Yet, both works only consider pitch-catch setups.

Outside the field of Non-Destructive Evaluation (NDE), the closest problem to on-plate ultrasonic mapping is acoustic room reconstruction. In this subfield of acoustics and signal processing, the objective is to reconstruct the shape of a room based on acoustic echoes on the walls. The design of robust methods for echo detection, and subsequent mapping, is explored in works like [79] for a single emitter/receiver, or [80] for a sensor array. Such methods, however, are only based on the reception of the first echo, whereas the other reflections also contain crucial information on the map. [69] elaborates a strategy to recover the geometry of a room with an arbitrary polygonal shape. Yet, it is only tested in simulation.

4.1.2 Simultaneous Localization and Mapping

Performing on-plate simultaneous localization and mapping by relying on omnidirectional UGWs is a challenging task. First, because SLAM is inherently a high-dimensional problem, where both the trajectory and map need to be jointly inferred. Second, due to the specificity of the ultrasonic measurements, which contain numerous (and potentially overlapping) echoes due to the multiple reflections on the plate edges. The consequence is that it is very challenging to

recover individual wave-packets from the mixture data [87]. Lastly, because of the echo-map feature association problem that arises when relying on the detection of individual echoes. In the literature, most UGWs techniques rely only on the incident wave packet [43, 86], especially for defect detection due to the difficulty of leveraging the multiple echoes. For on-plate localization and mapping purposes, however, the use of such reflections is essential, as they all provide range-only information to the edges.

In the literature, simultaneous localization and mapping based on the reflections of omnidirectional acoustic waves has been the subject of various works. [52, 53] address the problem of simultaneous room shape reconstruction and localization of the observation positions through batch optimization. In [51], a recursive Bayesian approach to echo SLAM is elaborated. Yet, all the mentioned works are only tested in simulation, and do not address explicitly the ambiguous matching of the echo measurements to room walls

Besides, several works have proposed acoustic SLAM frameworks that can work in practice, despite noisy measurements. [71] and [84] build a practical room reconstruction system running on a smart-phone. The latter would be used to emit sound pulses that would reflect from the room walls. In [71], both the localization and the wall estimation problems are considered, but using only the first echo contained in the measurements.

Overall, the development of a practical acoustic SLAM system that can leverage omnidirectional and multiple echoes is not established. More specifically, the data association problem in this context has not been thoroughly investigated, whereas it has always been a critical step, as described in reference works such as [66]. Besides, its efficiency may be key to accurate SLAM results. In the following, we present a first proof-of-concept method to achieve on-plate SLAM which has been elaborated before the start of this PhD thesis, and which attempts to explicitly solve the data association problem. It is based on FastSLAM, which has not been used previously in the context of acoustic reconstruction and localization problem, to the best of our knowledge. Due to the limitations in accuracy and robustness of this proof-of-concept approach, we next propose a novel development of FastSLAM that integrates beamforming maps and wave propagation models from the Rayleigh-Lamb theory. It has the particularity of not requiring explicit echo detection and association, while providing more accurate SLAM results.

4.1.3 Introduction to FastSLAM

FastSLAM [64] is an efficient approach to solve the online-SLAM problem. Let's consider that its resolution requires some correspondence variables $\mathbf{c}_{0:t}$ that are used to relate the observations to the map features, and which are assumed available in a first step. Based on the following decomposition of the target posterior over the map and trajectory:

$$p(\mathbf{x}_{0:t}, \mathbf{M} | \mathbf{u}_{0:t-1}, \mathbf{z}_{0:t}, \mathbf{c}_{0:t}) = p(\mathbf{x}_{0:t} | \mathbf{u}_{0:t-1}, \mathbf{z}_{0:t}) \prod_{k=1}^K p(\mathbf{m}_k | \mathbf{x}_{0:t}, \mathbf{z}_{0:t}, \mathbf{c}_{0:t}),$$

FastSLAM is implemented as a particle filter, where each particle contains a trajectory hypothesis and a collection of the map features estimates conditional on this trajectory. At time t , the particle set consists of N particle samples:

$$\mathcal{P}_t = \left\{ \mathcal{P}_t^{(n)} \right\}_{n=1..N} = \left\{ \mathbf{x}_{0:t}^{(n)}, \left\{ p(\mathbf{m}_k^{(n)} | \mathbf{x}_{0:t}^{(n)}, \mathbf{u}_{0:t-1}, \mathbf{z}_{0:t}, \mathbf{c}_{0:t}^{(n)}) \right\}_{k=1..K^{(n)}} \right\}_{n=1..N}.$$

One may note that this formulation allows for the number of map features to be estimated as well, as their number $K^{(n)}$ may vary with the particle. In typical uses of FastSLAM, the posteriors

over the map features can be recovered through a recursive mapping strategy such as Kalman filtering when the map is a collection of landmarks [64], or a binary filter when it is represented by an occupancy grid [94]. This results in an efficient implementation of the SLAM algorithm, as the trajectory hypothesis are sampled using the dynamic transition model, and the individual map features are determined independently using recursive methods.

In addition, FastSLAM allows handling uncertain data association between observations and map features by maintaining multiple particles with various association hypotheses. This can be important for the problem at hand as, when retrieving individual echoes from raw ultrasonic measurements, one may not know on which edge the reflection originates from. The use of a (naive) direct search approach for estimating the correspondence variables may be bound to failure, due to the paramount number of association possibilities for an entire set of measurements. To sidestep this issue, a more reasonable strategy is to affect the most-likely data association on a per-particle basis [94]. In addition, data association can be solved recursively to ensure a computationally efficient implementation to achieve online SLAM. Put in mathematical terms, the association for the n -th particle at time step $t + 1$ is determined based on past association estimates for the particle and Maximum Likelihood (ML):

$$\mathbf{c}_{t+1}^{(n)} = \arg \max_{\mathbf{c}_{t+1}} p \left(\mathbf{z}_{t+1} | \mathbf{x}_{0:t+1}^{(n)}, \mathbf{u}_{0:t}, \mathbf{z}_{0:t}, \mathbf{c}_{t+1}, \mathbf{c}_{0:t}^{(n)} \right).$$

For illustration, in a practical scenario where range measurements to fixed anchors would be received without prior knowledge on the correspondence between them, this strategy would associate each measurement to the anchor whose position estimate is the most consistent. In our case-study, the proof-of-concept method will adopt the same association strategy.

A last difficulty that may arise during a practical implementation, is that one may not know how many features need to be recovered a priori. Hence, for each observation, one should determine whether a new map feature is being observed, or whether the feature is already tracked by the filter. In the literature, algorithmic methods are proposed to tackle this issue, especially in reference books like [94]. A pseudocode version of the so-called FastSLAM 1.0 algorithm is provided in Algorithm 1 to present how its generic implementation can be structured to handle both data association and map feature initialization and tracking. For clarity of the presentation, it is assumed that a single measurement z_t is provided at a time, but the algorithm can seamlessly be extended to the case when multiple measurements would be received simultaneously, as we will see in the next section.

4.2 Proof-of-concept FastSLAM method

What follows is a brief recall of the proof-of-concept method presented in [70] to solve ultrasonic SLAM. The objective is to use UGWs to recover both the trajectory of an emitter/receiver pair of transducers and the geometry of the plate, while moving the sensor device on the plate surface. To support fast and robust echo detection, an \mathcal{L}_1 -regularized least squares approach to signal-delay estimation with a sparse signal model is used. This approach can result in a very accurate time-delay estimation of multiple overlapping reflections to be performed, even in low SNR conditions. The next section will describe how these components combine to demonstrate the feasibility of using UGWs for on-plate localization and plate geometry inference. The procedures for map initialization, map management and line-feature echo association are specified in detail.

Algorithm 1: FastSLAM_1.0($\mathcal{P}_t, \mathbf{u}_t, z_{t+1}$)

Data: Particle set \mathcal{P}_t , odometry data \mathbf{u}_t and measurement z_{t+1} .**Result:** Particle set \mathcal{P}_{t+1} for the next time step $t + 1$.*// Iterate over the particles***for** $n = 1..N$ **do***// Sample a new pose* $\mathbf{x}_{t+1}^{(n)} \sim p(\mathbf{x}_{t+1} | \mathbf{x}_t^{(n)}, \mathbf{u}_t)$;*// Determine the maximum likelihood data association* $\mathbf{c}_{t+1}^{(n)} = \arg \max_{\mathbf{c}_{t+1}} p(z_{t+1} | \mathbf{x}_{0:t+1}^{(n)}, \mathbf{u}_{0:t}, z_{0:t}, \mathbf{c}_{t+1}, \mathbf{c}_{0:t}^{(n)})$. ;*// Calculate the particle weight* $w_{t+1}^{(n)} \leftarrow p(z_{t+1} | \mathbf{x}_{0:t+1}^{(n)}, \mathbf{u}_{0:t}, z_{0:t}, \mathbf{c}_{0:t+1}^{(n)})$;**if** $w_{t+1}^{(n)} < p_0$ **then***// Initialize a new map feature if the association likelihood is lower than a threshold*Initialize $\mathbf{m}_{K^{(n)}+1}^{(n)}$ based on the measurement z_{t+1} ; $K^{(n)} \leftarrow K^{(n)} + 1$ **else***// Update the map feature k^* that is identified as being observed.*Update $p(\mathbf{m}_{k^*}^{(n)} | \mathbf{x}_{0:t+1}^{(n)}, \mathbf{u}_{0:t}, z_{0:t+1}, \mathbf{c}_{0:t+1}^{(n)})$ using the mapping strategy**end**Discard every $\mathbf{m}_k^{(n)}$ that has not been sufficiently observed. ;Update the total number of map features $K^{(n)}$ accordingly.**end**Construct \mathcal{P}_{t+1} by sampling each $\mathcal{P}_{t+1}^{(n)}$ proportionally to their respective weight $w_{t+1}^{(n)}$.**return** \mathcal{P}_t .

4.2.1 Notations and assumptions

We are considering a mobile unit transporting an acoustic emitter-receiver pair on a metal plate. At the i -th, the emitter sends an acoustic pulse $s(t)$. Under mild assumptions of the acoustic properties of the ultrasonic probing system, we can model the received signal $z_i(t)$ as a linear system through convolution:

$$z_i(t) = g_i(t) * s(t) + n_i(t), \quad (4.1)$$

with $g_i(t)$ the environment impulse response at \mathbf{x}_i and $z_i(t)$ which contains information about reflections on the plate edges, while n_i is an additive noise term. We still assume that the plate is a homogeneous material (steel, aluminum, ...) of a constant thickness. We assume the plate to be a convex polygon, but we do not assume it to be rectangular, even though this is to be expected in any industrial case. We will also consider that the ultrasonic echoes picked up by the receiver are first-order reflections. Any other reflections may lead to secondary echoes after bouncing from several edges, but these are neglected. It will also be assumed that odometry is known with good precision. More specifically, accurate heading is assumed to be available, as it can be provided by an accelerometer during a realistic scenario, on a nearly vertical structure.

4.2.2 Echo detection

One of the key enabling technologies for on-plate SLAM is the robust and accurate detection of ultrasonic echoes in the metallic plate. The presented approach applies a super-resolution technique based on \mathcal{L}_1 -regularized least-squares and the concept of sparsity to solve this estimation problem. The underlying hypothesis is that most of the metal plate does not reflect the ultrasonic waves, and that only discontinuities in the medium cause reflections to occur. In that case, the environment response $g_i(t)$ from Eq. (4.1) has only a limited number of non-zero components and can be considered sparse. Contrary to the signal model presented in the previous chapter that considers propagation models derived from the Rayleigh-Lamb equations, it will only be assumed that the measurements contain delayed copies of the emitted signal.

After discretization of the problem, these considerations allow us to cast the following \mathcal{L}_1 -regularized least squares problem for extracting echoes:

$$\hat{\mathbf{g}}_i = \arg \min_{\mathbf{g}} \|\mathbf{D} \cdot \mathbf{g} - \mathbf{z}_i\|^2 + \lambda \cdot \|\mathbf{g}\|_1,$$

where \mathbf{g} is the discretized vector-representation of the impulse response $g(t)$ of size $L \times 1$, \mathbf{z}_i the discretized received acoustic signal of size $L \times 1$, where L is the number of time-samples used in the discretization process. The notation $\|\cdot\|_1$ denotes the \mathcal{L}_1 -norm, with the well-known fact that it promotes sparsity of the least-squares solution [95]. The matrix \mathbf{D} is called a dictionary matrix of size $L \times L$, and contains time-delayed copies of the emitted echo signal \mathbf{s} :

$$\mathbf{D} = \begin{bmatrix} \mathbf{s}(1) & 0 & \dots & 0 \\ \mathbf{s}(2) & \mathbf{s}(1) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{s}(L) & \mathbf{s}(L-1) & \dots & \mathbf{s}(1) \end{bmatrix}$$

This minimization problem is a convex problem which can easily be solved using open-source toolboxes such as the Matlab CVX toolbox [34]. Based on the estimated sparse impulse response $\hat{\mathbf{g}}_i$, the times of arrival $\{\tau_{i,j}\}_{j=1..N_i}$ can be determined as the time delays associated with the non-zeros components of $\hat{\mathbf{g}}_i$. Assuming a known wave velocity v , the subsequent range measurements

$\{r_{i,j}\}_{j=1..N_i}$ are obtained using the formula:

$$r_{i,j} = \frac{\tau_{i,j}}{2v},$$

which is consistent with the back-and-forth travel of the wave. Due to noise and imperfect signal modelling, we will assume that not all the edges are detected all the time, and that spurious detection may be included in the list.

4.2.3 From echoes to edge hypotheses

Out of the detected ranges, a method is required to initialize plate edge estimates. This is necessary for initialization of the map and line generation using measurements that are not explained by the current map estimate. A single echo $r_{i,j}$ provides information about the distance to an object. As will be described later, this information is useful to refine the parameters of the edge which reflected this echo. However, this is too ambiguous to initialize a new edge: it could be any tangent to the circle centered on the current pose with radius $r_{i,j}$.

While three echoes are necessary to determine a line without ambiguity when the acquisition points are not aligned, relying on 2 echoes appears as a better option. When taken at different known positions, they are sufficient to define two line hypotheses: there are only two lines tangent to the two circles defined by these radii and centered on the known poses [51]. In addition, relying on 3 echoes would involve to consider numerous range triplets out of 3 sets of ranges. Also, this may not be necessary when moving in a direction normal to the edge, as two measurements are sufficient for a unique solution. Both cases are illustrated in Fig. 4.1. Overall, we only consider potential couples of ranges obtained at two different positions for initialization of a line.




Figure 4.1: Line hypotheses generated from two measurements at position P_{t-1} and P_t . Left: the robot moves away from the edge and generate a single hypothesis, right: generic case leading to two edge hypotheses.

Because of the combinatorial generation of edge hypotheses, a method is required to identify inconsistent edges and remove them. To this end, the hypothesis that only first-order reflections are reaching the receiver can be leveraged: with the assumption of a convex plate, if we consider two edges L_1 and L_2 observed from pose P and P_i , the projection of P on L_i , then L_2 cannot intersect the segment $[PP_1]$ and L_1 cannot intersect the segment $[PP_2]$. If one of these conditions is true, then both edges cannot exist simultaneously on a convex plate because observing one edge would require the UGWs to cross the other one, which is not possible. In the following, this test will be denoted as the consistency check between two edge hypotheses.

4.2.4 FastSLAM integration

With the aim to use FastSLAM for our localization and plate geometry inference problem, this section defines the particle state, the chosen map representation, the particle initialization strategy, and how the internal maps are updated based on the retrieved echoes.

FastSLAM state

Because of the assumption of a known orientation, we only consider the sensor coordinates (x, y) on the plate as a system state. A map is represented by a list of infinite lines, where each line is represented by an angle and an offset. We use a coordinate pair (α, b) to define the equation of a line:

$$x \cdot \cos \alpha + y \cdot \sin \alpha + b = 0$$

in a frame that is fixed with respect to the plate. A particle $\mathcal{P}_t^{(n)}$ is then described as a trajectory hypothesis paired with an estimated map:

$$\mathcal{P}_t^{(n)} = \left\{ (x_i^{(n)}, y_i^{(n)})_{i=0..t}, (\alpha_k^{(n)}, b_k^{(n)})_{k=1..K^{(n)}} \right\}$$

FastSLAM initialization

At initialization, we rely on edge hypotheses generated as explained earlier: after waiting for the second set of measurements, we estimate a number of edge hypotheses from which we extract maximally consistent sets. These sets are built using a dynamic programming approach inspired by JCBB/JCDA [66]. As a result, every particle is initialized by sampling around the zero position, and randomly selecting a consistent set of edges from the maximally consistent sets.

Particle evaluation

Given $\{r_{i,j}\}_{j=1..N_i}$ the set of echoes measured at the i -th scanning position, we evaluate a particle based on its ability to explain the measurements. For a line $L_k^{(n)} = (\alpha_k^{(n)}, b_k^{(n)})$ in particle $\mathcal{P}_t^{(n)}$, the expected measured range is:

$$d_{i,k}^{(n)} = |x_i^{(n)} \cdot \cos \alpha_k^{(n)} + y_i^{(n)} \cdot \sin \alpha_k^{(n)} + b_k^{(n)}|$$

From this range, the likelihood of measurement $r_{i,j}$ due to a reflection on the line $L_k^{(n)}$ is modeled with a Gaussian distribution centered on $d_{i,k}^{(n)}$, with a standard deviation consistent with the echo detection uncertainty. If the highest likelihood over all the lines is lower than a threshold, then the measurement is considered “unexplained” and allocated a low probability p_0 . Otherwise, the index of the line leading to the highest likelihood is recorded as $k_{i,j}^{\star(n)}$. For the purpose of map update, the list of unexplained measurements is also stored to create additional line hypotheses.

For a complete set of echoes, the evaluation of a particle will then be the product of the likelihood of the N_i likelihoods of the independent measurements:

$$\mathcal{L}(\mathcal{P}_t^{(n)}) = \prod_{j=1}^{N_i} \max \left(p_0, \max_k \left(p \left(r_{i,j} \mid \mathcal{P}_t^{(n)}, L_k^{(n)} \right) \right) \right).$$

This particle evaluation strategy effectively results in maximum likelihood data association, as each range measurement is associated with the line that best explains the observation.

Map update

The map update stage has three purposes, first for “explained” measurements, the associated line needs to be updated to account for the new piece of information. Second, the “unexplained” measurements are used to create additional line hypotheses. Third, consistency checks are used to eliminate line hypotheses which are no longer supported by the observations.

Line update: Given a range measurement $r_{i,j}$ and its associated line in particle $\mathcal{P}_t^{(n)}$, $L_{k^*}^{(n)}$, one needs to update the corresponding $(\alpha_{k^*}^{(n)}, b_{k^*}^{(n)})$. It should be noted that, for notation simplicity, we refer to the associated line index as k^* instead of $k_{i,j}^{*(n)}$ which makes explicit the time, particle and measurement indexes. Besides, we keep a record, on a per-particle basis, of the pairs of time step indexes and ranges associated with this line in a set $\mathbf{E}_{k^*}^{(n)}$. The line parameters update is then achieved by minimization of the following cost function:

$$C_{k^*}^{(n)}(\alpha, b) = \sum_{(i,j) \in \mathbf{E}_{k^*}^{(n)}} \left[(x_i \cdot \cos \alpha + y_i \sin \alpha + b)^2 - r_{i,j}^2 \right]^2,$$

New lines: For unexplained measurements, we need to combine them with the measurements from the previous pose to create edge hypotheses (as described previously). Every pair made of one previous measurement and a new unexplained measurement is used to generate one or two edge hypotheses. All the previous measurements are used for additional robustness against incorrectly associated measurements at the previous time step. All these hypotheses are added to the line set of the current particles.

Map clean-up: In a final stage of the map update, all the lines in a particle line set are checked for pairwise consistency. If two lines are deemed inconsistent and one of them has been observed (i.e. associated with a measurement) significantly more than the other since its creation, then the least observed line is marked for deletion. In our implementation, this criterion is defined as being observed two additional times. In this stage, we also mark for deletion lines that were not re-observed enough since their creation. For instance, we delete a line that is more than 10 step-old but has been observed less than 3 times. The simple delayed outlier removal described in this paragraph is what makes possible the generation of many edge hypotheses from the set of detected echoes.

4.2.5 Overview on the SLAM results

This section provides an overview of the SLAM results obtained with the proof-of-concept approach. We use the experimental data described in the previous chapter that was acquired on a $600 \times 450 \times 6$ mm aluminum plate, in a laboratory environment. We use an emitter/receiver pair of acoustic sensors that are moved, manually, to the predefined acquisition positions which are carefully recorded for ground truth positioning and plate boundary location.

Fig. 4.2 illustrates the echo detection process using an experimental measurement. Panel a) shows the emitted signal, which is a 2-cycle burst at 100kHz. The reflected signal from the plate structure can be seen in panel b). It shows the pickup of the emission at the beginning of the signal, and a reflection due to coupling mismatches. Then, a series of echoes is apparent. Using the \mathcal{L}_1 -based approach, the impulse response can be reconstructed (see panel b, orange trace). The method reconstructs the major echo components of the impulse response, while ensuring that the solution is sparse. Panel c) shows the reconstruction error. The fact that this residual is not zero can be explained by the fact that the signal model is not sufficiently appropriate to

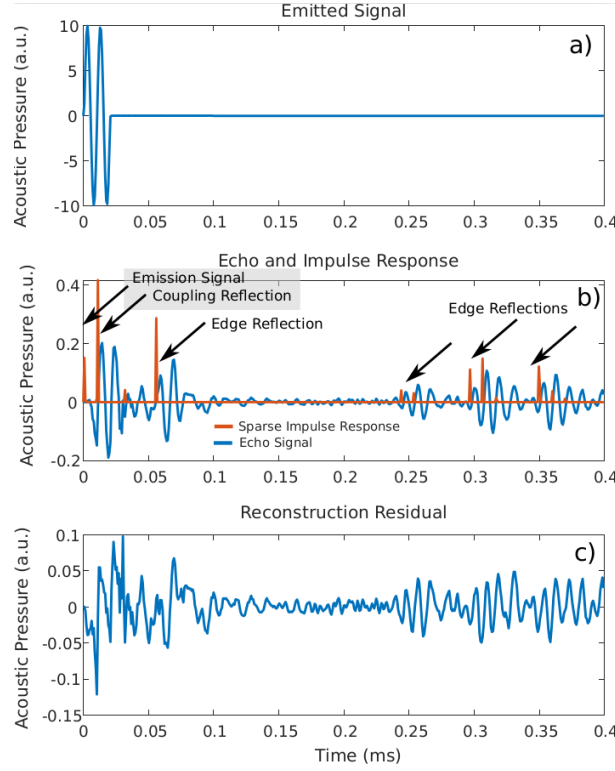


Figure 4.2: Illustration of the echo detection process. Panel a) shows the emitted signal. Panel b) shows the echo signal (blue trace) and the reconstructed impulse response (orange trace). Panel c) shows the reconstruction residual.




Figure 4.3: Evolution of the map representation for the highest-ranked particle at step 1, 36, 124 and 239. The purple frame represents the outline of the true plate. The red dot is the current estimated sensor pose, and the red line is the history of its estimated trajectory. The darkness of the line is proportional to the number of times they have been observed.

	Angle error [rad]	Offset error [cm]
Step 50	-0.029 ± 0.154	-6.1 ± 19.9
Step 239	-0.007 ± 0.091	-1.5 ± 12.6

Table 4.1: Average errors and standard deviations on the lines parameters estimation after 50 and 239 steps. They are evaluated using 100 repetitions.

describe the measurements. However, as the remaining experiments show, this crude approach to the echo detection process may be sufficient for the plate-estimation algorithm.

Figure 4.3 shows one anecdotal example of a mapping run on the plate by displaying the highest-ranked particle (out of $N = 16$ in this case), its line set and trajectory. In this scenario, the transducers follow a simulated lawn-mowing trajectory. The reconstructed path is clearly visible, and the stability of the estimate can be observed when performing a second sweep of the path (iteration 239). The convergence of the estimated lines to the true plate outline can also be observed, with a final error in the order of 2cm on the line offsets b .

To numerically evaluate the performance, we provide, in Table 4.1, the mapping precision and repeatability of this SLAM approach through 100 repetitions of the considered scenario. The table presents a summary of the resulting precision at different steps in the SLAM process. It is clear that the initial estimates, after observing only half of the plate, are still very uncertain. Although there is little bias in the estimation during the last step, the standard deviation values are relatively important, even after a second sweep of the plate. Overall, this proof-of-concept approach demonstrates the feasibility of using ultrasonic data for on-plate SLAM, as the ultrasonic measurements contain information to the distance to the plate edges. Still, one may seek to improve its accuracy and robustness due to the high variability of the reconstruction results.

4.3 A FastSLAM approach integrating beamforming maps

We now present a novel FastSLAM approach to achieve on-plate SLAM by relying on ultrasonic reflections on the plate boundaries. In the previous experiments, ultrasonic measurements acquired on metal plates have proven to yield sufficient information to provide both localization and mapping capabilities. It is based on \mathcal{L}_1 -regularized least squares to retrieve the multiple echoes within one measurement. However, the dispersive nature of the waves was not taken into account, and the relative complexity of the algorithm due to explicit echo-line association and the map management strategy may jeopardize its robustness and accuracy.

Contrary to the previous method, our new approach takes into account the dispersive nature of guided waves in metal plates. Indeed, we rely on a wave propagation model to determine, through correlation with acoustic data, the likelihood that a reflection occurred to a wide range of distances to the transducers. These correlation measurements are then leveraged to construct beamforming maps that focus the measurement energy in the 2D-line space. Based on the latter, we solve the mapping problem through plate edges estimation for every particle, in a FastSLAM fashion. In this new setup, the resolution of the difficult echo association problem is no longer required. Yet, for simplicity, this approach is restricted to recovering rectangular shapes. It will be demonstrated, with real acoustic measurements obtained on different metal plates, that such a framework achieves more accurate results than the proof-of-concept method, while the complexity of the algorithm is sensibly reduced.

4.3.1 Plate boundary mapping with delay-and-sum beamforming

We will be interested in mapping the plate boundaries using the ultrasonic reflections, as in the previous approach. As the overall SLAM method will be integrated within FastSLAM, we can assume first a known trajectory, as different trajectory hypotheses will be maintained by the particle set. Here, we propose a mapping approach based on the correlation signals presented in Chapter 3, with the aim to recover a rectangular plate geometry estimate. To compute these signals, we rely on the acoustic model presented in Chapter 2:

$$\hat{g}(r, \omega) \approx e^{-jk(\omega)r} / \sqrt{k(\omega)r}, \quad (4.2)$$

where $k(\omega)$ is the wavenumber of the major acoustic mode (A0). The correlation measurement is determined with:

$$z'_i(r) = \frac{\langle z_i(t), \hat{z}(r, t) \rangle}{\sqrt{\langle z_i(t), z_i(t) \rangle \langle \hat{z}(r, t), \hat{z}(r, t) \rangle}},$$

and its envelope is retrieved with:

$$e_i(r) = |z'_i(r) + j\mathcal{H}(z'_i)(r)|.$$

The determination of e_i serves the purpose of echo detection, as the local maxima provide the most-likely ranges to the plate edges, when considering echoes that are first-order ones. This principle has already been illustrated in Chapter 3.

Similarly to the proof-of-concept method, the map is represented by a set of lines: $\mathbf{M} = \{r_l, \alpha_l\}_{l=1\dots 4}$, where the parameters (r_l, α_l) define the line equation in the 2D plane with:

$$x \cdot \cos \alpha_l + y \cdot \sin \alpha_l - r_l = 0$$

in a non-mobile frame with respect to the plate. In this thesis, we will often use the following convention: the origin of the reference frame will be taken as the initial robot position, and the x -axis will be aligned with the initial robot heading, as illustrated by Fig. 4.4. Moreover, we limit our case-study to rectangular shapes. Hence, the possible maps possess only four lines forming a rectangle altogether. This is not a major limitation for inspection on ship hulls and storage tanks, as their structure is almost entirely made of rectangular panels. Still, the next chapter of this thesis will focus on more challenging mapping problems. Especially, an extension to the arbitrary polygonal shape will be investigated.

It is noteworthy that one single range measurement is not sufficient for recovering plate boundaries without ambiguities, as all the lines tangent to the circle which radius equal to the detected range, and centered at the sensors' position, may equally account for the correlation measurement. Consequently, measurements e_i acquired at different positions must be combined. Let's assume a hypothetical robot trajectory $\mathbf{x}_{0:t}$. We aim at estimating the map \mathbf{M} , which means inferring a probability density function $p(\mathbf{M}|\mathbf{x}_{0:t}, z_{0:t})$. A first solution would consist in assessing, for each map in the 8-D domain, the correlation between the observations and the predicted data based on the image source model. However, such an approach would be far too cumbersome for a real-time application. Instead, we rely on a Delay-and-Sum (DAS) beamforming approach. We determine a beamforming map that attributes, to every line parameters (r, α) , the likelihood of the line existence given the observations with:

$$\mathcal{L}_t(r, \alpha) = \sum_{i=0}^t e_i(d_i(r, \alpha)) = \sum_{i=0}^t e_i(|x_i \cdot \cos \alpha + y_i \cdot \sin \alpha - r|).$$




Figure 4.4: Representation of lines using (r, α) coordinates. As a convention, the reference frame is taken as the initial robot position, and the x axis is aligned with the initial robot heading.

where $d_i(r, \alpha) = |x_i \cdot \cos \alpha + y_i \cdot \sin \alpha - r|$ is the distance between the robot during time-step i and the hypothetical line being considered. In the equation, all the correlation values add up constructively along all the observations if an edge is indeed present. Also, it can be noted that only first-order reflections are taken into account, as we reason on individual lines. One may consider that higher order reflections, which we consider as "noise", are less likely to account for high correlation amplitudes because of wave scattering after each additional reflection. Indeed, the latter causes loss of energy to the wave packet. One major advantage of this approach is that $\mathcal{L}_t(r, \alpha)$ can be computed recursively when an additional measurement e_t is made available, as $\mathcal{L}_t(r, \alpha) = \mathcal{L}_{t-1}(r, \alpha) + e_t(|x_t \cdot \cos \alpha + y_t \cdot \sin \alpha - r|)$. This is beneficial for a robotic task meant to be performed in real time, as a map estimate is available at any time, and the computational load of one update is low. Finally, to retrieve a maximum likelihood estimate of the map, we solve the following optimization problem:

$$\hat{\mathbf{M}}_t = \arg \max_{\mathbf{M}} \mathcal{L}_t(\mathbf{M}) \triangleq \arg \max_{\mathbf{M}} \sum_{(r_l, \alpha_l) \in \mathbf{M}} \mathcal{L}_t(r_l, \alpha_l), \quad (4.3)$$

where \mathbf{M} is restricted to be a rectangle. It can be solved efficiently by taking that constraint into account. First, one can determine the most likely line:

$$(\hat{r}_1, \hat{\alpha}_1) = \arg \max_{r, \alpha} \mathcal{L}_t(r, \alpha).$$

Based on a discretized version of \mathcal{L}_t , the above optimization can be solved seamlessly with a grid search. Next, it is possible to rely on the assumption that the retrieved line provides the most reliable estimation of the plate orientation with respect to the robot. Therefore, the determination of the other lines for $l = 2, 3, 4$ reduces to solving simple and independent one-dimensional optimization problems:

$$\hat{\alpha}_l = \hat{\alpha}_1 + \frac{\pi(l-1)}{2} ; \quad \hat{r}_l = \arg \max_r \mathcal{L}_t(r, \hat{\alpha}_l).$$

Overall, this mapping approach is relatively simple. Not only it presents little computational complexity and can be deployed seamlessly online, as we will show, but also, no explicit echo detection and association are required.

4.3.2 Particle evaluation and FastSLAM algorithm

Similarly to the proof-of-concept method, we rely on a FastSLAM approach. We use a particle filter in the localization space, where each particle holds a hypothesis on the map which is inferred from the particle trajectory and the measurements. In our approach, a set with N particles has the following form during measurement step t :

$$\mathcal{P}_t = \left\{ \mathbf{x}_{0:t}^{(n)} = [x_i^{(n)}, y_i^{(n)}, \theta_i^{(n)}]_{i=0..t}, \mathcal{L}_t^{(n)} \right\}_{n=1..N}$$

where $\mathbf{x}_{0:t}^{(n)}$ represents the n -th particle belief on the robot trajectory augmented with its heading over time steps $i = 0..t$, and $\mathcal{L}_t^{(n)}$ its beamforming map which depends on the particle trajectory hypothesis. Based on $\mathcal{L}_t^{(n)}$ and the resolution of Eq. (4.3), each particle holds a hypothesis $\mathbf{M}_t^{(n)}$ on the map. Besides, each particle is provided with a weight indicating how the particle belief accounts for the measurements. To define it, we rely on the current envelope of the correlation measurement, and assess the likelihoods of the map edges retrieved from $\mathcal{L}_t^{(n)}$ and the current robot position belief:

$$w_t^{(n)} = \eta \cdot \exp \left\{ \beta \sum_{(r_l, \alpha_l) \in \mathbf{M}_t^{(n)}} e_t \left(d_t^{(n)}(r_l, \alpha_l) \right) \right\} \quad (4.4)$$

where η is the normalization factor and β a positive parameter. It enables to fix the confidence in the correlation measurements and shall be tuned so that the resulting weight distribution is consistent with the motion and observation noises. One will notice that Eq. (4.4) is the same observation model as the one used in Chapter 3, except that the plate boundaries location are now estimated. The calculated weights are used to sample, with replacement, the particles after each time step. Besides, one may note that we are not determining a full posterior over the map, thus, in Eq. (4.4), the uncertainty on the lines' retrieval from the beamforming maps is not considered. Altogether, our implementation of FastSLAM is given in Algorithm 2.

Algorithm 2: FastSLAM($\mathcal{P}_{t-1}, \mathbf{u}_{t-1}, e_t$)

Data: Particle set \mathcal{P}_{t-1} , odometry data \mathbf{u}_{t-1} and correlation measurement $e_t(r)$.

Result: Particle set \mathcal{P}_t for the current time step t .

if $t=0$ **then**

 Initialize the particle filter with $\mathcal{P}_0 = \{[x_0, y_0, \theta_0], \text{null-fuction}\}_{n=1..N}$

else

for $n = 1..N$ **do**

$\mathbf{x}_t^{(n)} \sim p \left(\mathbf{x}_t | \mathbf{x}_{t-1}^{(n)}, \mathbf{u}_{t-1} \right)$;

$\mathcal{L}_t^{(n)}(r, \alpha) = \mathcal{L}_{t-1}^{(n)}(r, \alpha) + e_t(|x_t^{(n)} \cos \alpha + y_t^{(n)} \sin \alpha - r|), \forall (r, \alpha)$;

$\mathbf{M}_t^{(n)} = \arg \max_{\mathbf{M}} \mathcal{L}_t^{(n)}(\mathbf{M})$;

$w_t^{(n)} \propto \exp \left\{ \beta \sum_{(r_l, \alpha_l) \in \mathbf{M}_t^{(n)}} e_t \left(d_t^{(n)}(r_l, \alpha_l) \right) \right\}$

end

 Construct \mathcal{P}_t by sampling each particle proportionally to their respective weight.

end

return \mathcal{P}_t .




Figure 4.5: Acquisition positions on the steel plate.

4.4 Experimental results

In this section, we test our FastSLAM approach on experimental data. We detail the experimental setups and show the results in terms of localization and mapping accuracy. We will also compare the localization and mapping performance with that of the proof-of-concept method. To illustrate the feasibility of ultrasonic SLAM with the proposed method and using a real robotic platform, the approach is tested on data acquired with a magnetic crawler.

4.4.1 Experimental setups

To assess the efficiency of our procedure, we will consider two scenarios. In Scenario 1, we use the acoustic data acquired on the 600x450x6mm aluminum plate, which has already been presented. In the second scenario, we consider a larger rectangular metal plate which has dimensions 1700x1000x6 mm, and is made of steel. The acquisition process is globally the same to collect the data on this second plate: the transducer pair is moved by hand on the vertices of a regular grid. At every position, 10 measurements of the ultrasonic response are averaged to improve the signal quality. This operation is not critical in a laboratory environment, but it may be necessary in outdoor conditions, where the inspection robot shall operate, to alleviate the effect of external disturbances. The acquisition positions, which are carefully recorded for ground-truth positioning, are provided in Fig. 4.5.

In total, 108 measurements are collected on the aluminum plate (plate 1), while this number increases to 117 for the steel plate (plate 2). We use two tonebursts of a sinusoidal wave at 100 kHz as the excitation. Moreover, the direct incident signal is smoothly removed from the data using a sigmoid window, as it does not correspond to a reflection on an edge.

For each plate, we determine a wave propagation model as in Eq. (4.2). More specifically, we solve the Rayleigh-Lamb equations and assume that the A0 Lamb wave mode is predominant in the two scenarios. The model parameters ($d = 6$ mm, $c_L = 6420$ m/s and $c_T = 3040$ m/s) are used for Scenario 1, and have been experimentally validated in Chapter 3. We use the model parameters: $d = 6$ mm, $c_L = 5880$ m/s and $c_T = 3250$ m/s for the second scenario, which have also been experimentally validated by comparison between the theoretical and experimental dispersion curves.

Regarding the particle filter parameters, we use $N = 20$ particles. To simulate a sweep of a plate by a robotic platform, a sequence of measurements is selected from the database and is

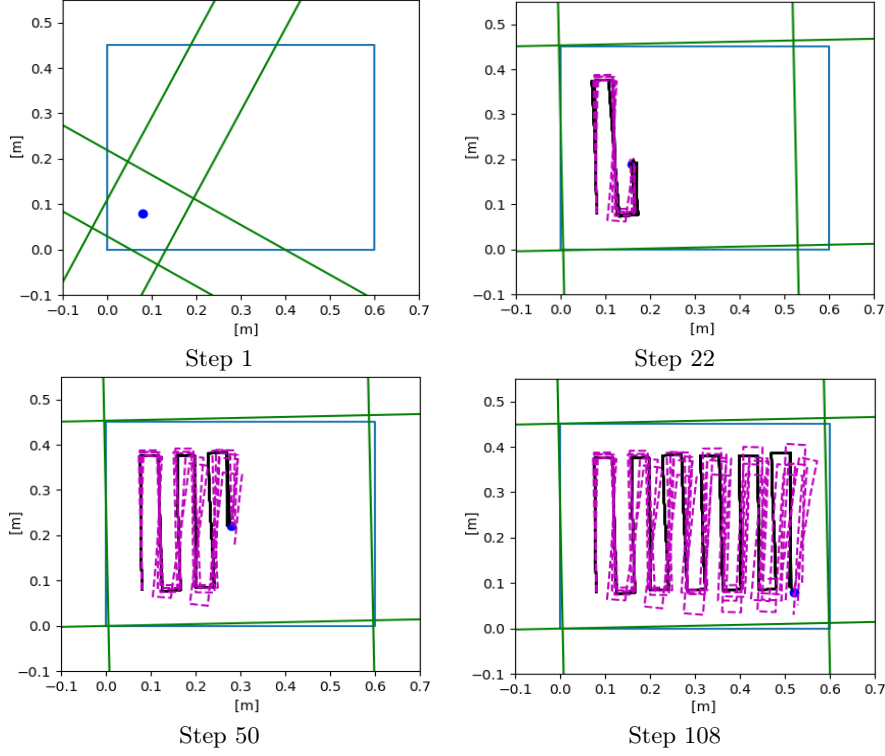


Figure 4.6: Trajectories estimated by all the particles (black lines), dead-reckoning trajectories (dash magenta lines) and map retrieved by the most likely particle (green lines) during Steps 1, 22, 50 and 108 for a lawn-mower path on plate 1 (zoom for details). The true outline of the plate and true sensor positions correspond to the blue rectangle and blue dot respectively.

presented to the SLAM framework, with the theoretic displacement between grid cells used as odometry. Also, we add Gaussian noise to the odometry data: $\bar{\Delta}r \sim \mathcal{N}(\Delta r, (10^{-2}\Delta r + \Delta r_0)^2)$ and $\bar{\Delta}\theta \sim \mathcal{N}(\Delta\theta, (10^{-2}\Delta\theta + \Delta\theta_0)^2)$ with $\Delta r_0 = 10^{-3}$ m and $\Delta\theta_0 = 10^{-2}$ rad to simulate odometry uncertainty which is considered to be moderate due to the robot magnetic adherence and embedded accelerometers that may be used to provide precise heading on a nearly-vertical structure, in a realistic scenario.

4.4.2 Localization and mapping results

We run our FastSLAM algorithm using the data of plate 1, and simulate a lawn-mower path. Although the results are generated off-line, our method can run online on a real robotic platform. Indeed, as the beamforming maps of size $Z \times Z$ are updated incrementally, the complexity of one FastSLAM iteration with N particles is $\mathcal{O}(N \times Z^2)$, which leads to a computational time no higher than 100 milliseconds per iteration in our setup, with $Z = 300$ and $N = 20$. Consequently, the FastSLAM updates can reasonably be preformed at a rate of 5 to up to 10 Hz. The particle filter is initialized with $[x_0, y_0, \theta_0] = [0, 0, 0]$.

In Fig.4.6, we show the particles' belief on the sensors trajectory during measurement steps 1, 22, 50 and 108. We also represent the map retrieved by the particle with the highest weight and several dead-reckoning trajectories obtained using noisy odometry input only. During Step 1, the map is not correctly estimated. As only one measurement has been integrated, the distance to the closest edge can be recovered but, the orientation is essentially random. Rapidly, the three closer




Figure 4.7: Beamforming map for the particle with the highest weight during the final step. The (r, α) coordinates are expressed in the reference frame which origin is the initial position of the sensors, and the x axis is aligned with the initial heading. The rectangles indicate the edges retrieved with our method.

edges are recovered as shown during Step 22. However, the right edge is not yet well estimated as it is further away. During Step 50, the plate shape is fully recovered, and during the final step, both the estimation of the plate shape and trajectory are accurate. In contrast, the dead-reckoning trajectories present noticeable drift. This illustrates that, by relying on the acoustic data, the proposed approach can appropriately compensate for moderate odometry noise.

Fig.4.7 depicts the beamforming map for the most likely particle during the final step. We can see that the intensity peaks due to the edges can be easily identified, and our optimization method does not face difficulty to retrieve them.

To compare the proposed FastSLAM approach with the proof-of-concept one, we show, in Fig.4.8, the average localization and line parameters estimation errors calculated over 100 runs of each algorithm, and using the same acoustic data. The localization is determined by evaluating the distance between the x, y sensor position in the estimated plate frame, and the ground truth position in the true plate frame. We simulated 100 repetitions of the lawn-mower path for the sensors' trajectory. In the figure, we also represent the 10% and 90% quantiles with the aim to measure the repeatability of each approach. It can be observed that, with our new method, only a few tens of measurement steps are necessary to recover, in average, the range parameters of the lines with a precision of a few millimeters, and the plate orientation with a precision better than one degree. The localization result is also very precise as, after a quick convergence, the position errors remain in the order of a few millimeters despite the defects on the plate. Besides, the estimation is not subject to randomness, as the 10 % and 90% quantiles remain close to the average results. In comparison, the proof-of-concept method demonstrates poorer results. Indeed, not only are the average estimation errors higher, but also the variation of precision can be relatively significant between two runs. We have seen earlier that it required many more steps before more accurate results can be achieved. Also, it is sensitive to echo detection and association errors. Altogether, the results illustrate the improvement on localization and mapping that is achieved with the new method.

With the aim to assess the performance on a larger plate, we run our algorithm with the measurements obtained on plate 2, and simulate, again, a lawn-mower path. The results obtained

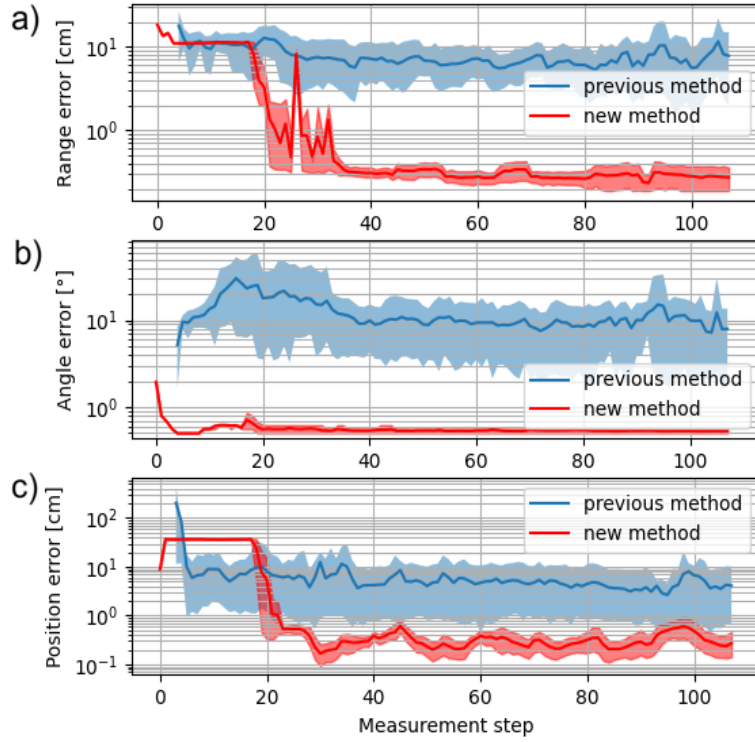


Figure 4.8: Localization and mapping results over 100 repetitions of a lawn-mower path on plate 1 for the previous and the new FastSLAM methods. a) Average estimation errors on the range parameter of the lines. b) Average estimation errors on the angle parameter. c) Average localization errors in the estimated plate frame. The 10% and 90% quantiles correspond to the upper and lower bounds of the coloured areas. The scales along the y-axis are logarithmic.




Figure 4.9: Localization and mapping results over 100 repetitions of a lawn-mower path on plate 2 for the new FastSLAM method. a) Average estimation errors on the range parameter of the lines. b) Average estimation error on the angle parameter. c) Average localization errors in the estimated plate frame. The 10% and 90% quantiles correspond to the upper and lower bounds of the colored areas. The scales along the y-axis are logarithmic.

Scenario	Range error [mm]	Angle error [degree]
Scenario 1	3.007 ± 0.098	0.234 ± 0.0004
Scenario 2	10.766 ± 22.921	0.206 ± 0.134

Table 4.2: Average estimation errors and standard deviations on the lines parameters obtained during the last measurement step for the two scenarios in consideration. The errors are evaluated using 100 repetitions.

over 100 runs are provided in Fig.4.9. On this plate, the echo detection employed by our previous method is not efficient, as it does not consider the wave dispersion effect, whereas the propagation distances are larger. This induces large misdetection rates and poor results. Hence, we display only the results of our new approach. Despite the slower convergence caused by the larger surface, and the slightly higher localization error, our method still provides precise estimates of the trajectory and plate geometry. This result indicates that our approach still works on surfaces sufficiently large to be used for realistic applications. The underlying prerequisites are a wave propagation model and filter parameters that conveniently fit the acoustic measurements and on-the-field noisy conditions. Naturally, one may also expect longer convergence times when the plate surface increases, as the echo detection is expected to be efficient mostly for short ranges as shown in Fig.4.9.a).

As a final evaluation, we determine the average mapping errors and standard deviations over 100 runs obtained during the final step for a lawn-mower path (Scenario 1) and a random walk (Scenario 2) on plate 1. Fig. 4.2 presents the results. It can be noticed that the overall results are relatively poorer for the random walk. This illustrates that the estimation accuracy also strongly depends on the robot path which shall be optimized for optimal reconstruction.

4.4.3 Experiment with a magnetic crawler

We evaluate our approach using a real robotic system acquiring ultrasonic data on a large plate to demonstrate its feasibility for a practical robotic inspection task. In our experimental setup, a magnetic crawler is moving on the surface of the 1700x1000x6mm steel plate, in an indoor laboratory environment. It is equipped with an onboard accelerometer to provide accurate heading measurements, and with a single transducer that is mounted in its head for simultaneous acoustic emission and reception.

This true pulse-echo setup (contrary to the pseudo-pulse echo setup with a pair of emitter/receiver used in the previous experiments) requires using an electric circuit for protecting the acquisition device during the high-voltage excitation. The prototype circuit designed to achieve that purpose is provided in Fig. 4.10. The principle is to protect the acquisition device during the excitation at high amplitude, otherwise, the acquisition device could be damaged. This is achieved by the use of Zener diodes which are placed in parallel with the acquisition device, as they can saturate the voltage level at their ends. The remaining of the excitation power is dissipated by the resistor. In this circuit, the commutation diodes placed in parallel, and in series with the generator, act as an interrupter: when the generator is no longer providing an excitation (*i.e* when the voltage V and intensity go to zero), it is "disconnected" by the diodes from the rest of the circuit to avoid short circuit of the echo measurements.

Despite the fact that echo signals should be preserved by the circuit, this setup leads to importance disturbance of the signals. Thus, care is taken in designing a sigmoid filter window and a band pass filter to obtain signals that can be used by our FastSLAM algorithm.

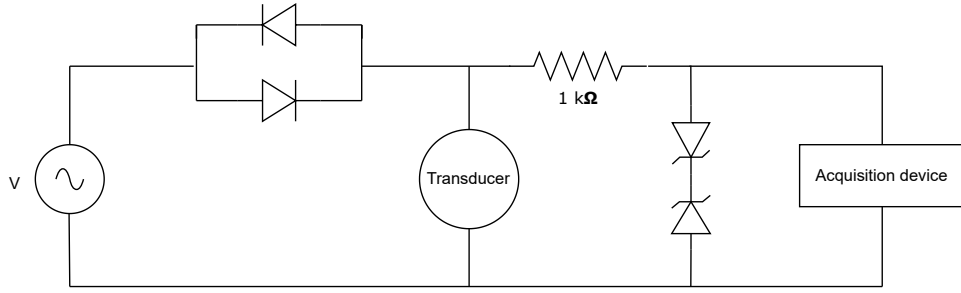



Figure 4.10: Schematic of designed electrical circuit for a true pulse-echo setup. The diodes in parallel are 1N4148 commutation diodes, while the diodes in series are Zener diodes.

To acquire the measurements, the robot is driven along a trajectory on the plate positioned nearly vertically, and the acquisition of one measurement is triggered each time the sensor has traveled 5 cm. Also, a pump is activated to continuously bring water through the tether at the interface between the sensor and the plate for sufficient coupling. Due to variations of the contact condition between the transducer and the plate surface, the varying coupling level maintained by the water supply and the electric circuit, the signal-to-noise ratio in the measurements is relatively poorer than when the acquisition is carefully performed by hand. We will see that, in these conditions, localization and mapping can still be achieved.


In Fig. 4.11, we show qualitative results obtained with our FastSLAM method during different steps of the considered scenario. An accurate localization system such as a laser total station could be used to precisely assess the localization error by tracking the position of the sensor. However, such a system was not available at the time of the experiment. Also, snapshots of the camera recording of the experiment are provided for each corresponding step. It can be seen on the figure that, during Step 1 which is the initialization step, a square estimate of the plate geometry is returned, as only a single measurement has been integrated. During Step 2, the robot has moved to the bottom part of the plate. Yet, the mapping error appears seemingly important, and the uncertainty on the robot position and orientation is large. The reason is that the rotational noise is probably overestimated in the transition model, and the map estimate, which is still erroneous, cannot be used for correction. This issue starts to resolve itself during Step 3, where the position of the particles in the reconstructed plate appears coherent with the reality. This remains the case until the final step. Overall, this illustrates the feasibility of using ultrasonic echoes in practice to achieve localization and mapping on a metal plate with a real robotic platform, and based on the proposed FastSLAM approach.




Step 1: Camera snapshot




Step 1: SLAM results




Step 2: Camera snapshot




Step 2: SLAM results



Step 3: Camera snapshot



Step 3: SLAM results



Step 4: Camera snapshot



Step 4: SLAM results

Figure 4.11: Qualitative SLAM results obtained with a magnetic crawler, and using the proposed FastSLAM approach, during different steps of the experiment.

4.5 Chapter summary

We have designed a new FastSLAM approach to achieve Simultaneous Localization and Mapping on metal plates by relying on the reflections of an ultrasonic excitation on the plate boundaries. Contrary to the proof-of-concept approach, this new method relies on a wave propagation model from the Rayleigh-Lamb theory and beamforming maps for estimating the plate edges location. Also, this method has the particularity of not requiring explicit echo detection and association, which are challenging problems, especially when considering ultrasonic data.

Experiments carried out on an undamaged and a damaged metal plate in a laboratory environment demonstrate that this new approach achieves better results in terms of localization and mapping accuracies and robustness, with less algorithmic complexity. More specifically, it was demonstrated that (sub-) centimeter precision can be achieved in localization and mapping in moderate odometry noise conditions. Yet, the proposed approach is restricted to recovering rectangular plate geometries. Finally, the approach has been validated using a real robotic platform, in a laboratory environment. Despite the low quality of the signals due to noise, the qualitative results illustrate the feasibility of ultrasonic-based on-plate SLAM with the proposed approach.

Chapter 5

Guided wave-based mapping with combined representations

The ability to reliably map acoustic scatterers using ultrasonic echoes is key to successful long-range robotic inspection. By relying on beamforming to locate the boundaries of a metal plate, the feasibility of on-plate ultrasonic SLAM has been demonstrated in the previous chapter. However, such an approach has some limitations: although it is relatively simple, this approach does not tackle a well-known issue of DAS beamforming, which is signal interference. Furthermore, the resulting point-estimate map lacks of uncertainty assessment. As a result, the confidence one can put in the map estimate may be unclear, especially in a realistic and noisy environment.

In this chapter, we introduce a more elaborate mapping framework based on combined map representations. More specifically, we seek to recover two different (grid and feature-based) maps of the same metal plate environment. We rely on the dense and probabilistic grid map for uncertainty assessment, and subsequently leverage the property of consistency between the two representations to mitigate interference. As the proposed framework relies on the detection of the closest edge reflector, probabilistic models are introduced to ensure robustness with respect to echo detection error. The approach is extensively evaluated in simulation against a more standard mapping method, and it is experimentally validated in a laboratory environment.

In the second part of this chapter, the designed mapping framework is integrated within a frontier exploration strategy with the aim to automatically drive an autonomous robot towards informative measurement locations. We determine that such a strategy may result in faster plate reconstruction and coverage compared to standard exploration strategies.

Last, as the reconstructed plate geometries were, so far, limited to rectangular shapes, we see how a partially reconstructed grid map can be used to extend feature-based plate boundary mapping to recover more arbitrary polygonal shapes from the ultrasonic echoes. The approach is extensively evaluated in simulation and in a real-world setting.

5.1 Motivation and background

The experimental results from the previous chapter have shown that delay-and-sum beamforming can be efficient for mapping a rectangular metal plate using UGWs. This method is relatively simple to implement, and it can seamlessly run online, within a FastSLAM framework. However, these benefits do not come without disadvantages and limitations. First, DAS beamforming is a non-adaptive method (*i.e* the way how different measurements are fused together for mapping does not depend on the data itself), and it does not tackle the detrimental effect of signal

interference (that will be further explained in the next section), whereas it is a well-known issue of DAS beamforming. Second, beamforming only provides point map estimates. As a result, due to the lack of uncertainty assessment and the time-inconsistent map estimates provided in the early steps, the approach is not suitable for mapping several metal plates on a large structure, which may require robust detection of plate boundary-crossing. Indeed, detection failures may result in mapping inaccuracies, as the present DAS approach requires that the measurement positions lie on the same metal panel. Besides, the lack of uncertainty assessment makes the approach unsuitable for active sensing, whereas it has been shown that the robot trajectory affects the estimation results. Last, this mapping method can only recover rectangular geometries. Although it may not be an issue for mapping metal panels on ship hulls and storage tanks, the extension to arbitrary shapes can be beneficial for application to more general structures, or even broader, to a wider range of acoustic SLAM problems.

We will see in this chapter that, surprisingly, the use of combined map representations may help tackle the three aforementioned limitations. We will aim at recovering grid-based and feature-based spatial representations jointly, while still limiting our case-study to mapping metal plate boundaries. In this context, the feature-based map will be represented by a set of two-dimensional lines, each line accounting for a potential boundary. The grid will consist of a dense representation of the plate environment that can account, all along the mapping process, both for areas identified as inside or outside the current metal plate, and areas whose state is still unknown, making it is suitable *e.g.* for detecting a change of plate, or for use in a later active-sensing strategy. The construction of the grid is based on the Occupancy Grid (OG) framework, widely used in robotics, that will be briefly recalled.

Maintaining two different representations of map estimates brings its own challenges. As they should represent the same environment, they *must* always remain consistent in some sense. Yet, this consistency property shall be leveraged: on the one hand, the grid map is more spatially informative than the feature-based map, which can help to filter out candidate plate edges that are no longer relevant, and mitigate the detrimental effect of interference. On the other hand, line features estimates, when associated with individual echoes, can be useful to estimate areas outside the plate for more robust boundary detection. In this chapter, we will attempt to establish a mathematical framework for building such joint maps. It is tested in simulation, with measurements acquired experimentally by hand, and with a magnetic crawler.

Besides, the past experiments have highlighted the strong dependency of the reconstruction results with the robot's trajectory. Hence, a legitimate question would be: how the robot should navigate to maximize the information gain on the map? Can it be achieved algorithmically? Based on the availability of the probabilistic grid map, we will see that a frontier exploration strategy appears to be a reasonable option, as the robot is steered towards unexplored areas while aiming for a maximum mapping coverage. Also, one may easily comprehend that such a strategy is relevant for robotic inspection as, after operation, all the structure surface must be inspected (or "explored").

Finally, to alleviate the need of the rectangular geometry assumption, we will see that leveraging a consistency constraint between the line feature-based map with a (partially reconstructed) grid map can help to recover arbitrary polygonal shapes, without prior knowledge on the number of boundaries to estimate. Although its practical benefit for our target application may be limited, this new approach may provide an elegant solution to the challenging geometry reconstruction problem using omnidirectional acoustic waves. Indeed, through minimization of a single cost function, an arbitrary shape could be recovered without any explicit echo detection and echo association, resulting in little algorithmic complexity and sensitivity to noise.

5.1.1 Related works

So far, the mapping methods presented earlier are exclusively feature-based, as they can recover metal plate boundaries as line features using DAS beamforming. Yet, grid-based map representations are widely used for robotic applications, especially for exploration [94, 33, 102], as, contrary to feature-based maps, they can explicitly account for explored and unexplored areas.

The estimation of grid maps using UGWs and a moving sensor has been investigated in related works such as [90, 25, 91]. However, the proposed methods rely on directional GWs, which necessitates constant alignment of the sensor directivity with the vector perpendicular to the boundary. In other words, these works rely on substantial prior knowledge on the map for robot navigation. In [62], omnidirectional GWs are used to autonomously map, in a pitch-catch setup via a frontier exploration strategy, the inside of a metal panel by relying on a grid representation. However, the multiple echoes within the signals are not exploited in the mapping procedure, and the presence of echo measurement errors is not considered.

In the robotics literature, several works have demonstrated the potential of combining grid and feature-based maps using conventional directional sensors (*e.g.* camera, lidar, sonar) [5, 68, 101]. One essential motivation behind this approach (along with the cumulation of both map sparsity and denseness) is that sparse map features can be used to integrate dependency between the measurements to build the grid map, whereas they are almost systematically assumed independent in the conventional OG framework. In the proposed framework, the consistency between the two maps is leveraged to mitigate the detrimental effect of interference. In addition to the use of DAS beamforming, our method is built on the detection of the *single* echo due to reflection on the closest edge. While the corresponding detection errors can be a source of mapping inaccuracies, our probabilistic formulation enables our approach to be robust to such. In the literature, however, such detection issues are often overlooked by recent works addressing the similar problem of room shape reconstruction from acoustic echoes in the air [51, 69].

5.1.2 Binary Bayesian filtering

The binary Bayesian filter is conventional for estimating occupancy grid maps. In this representation, each grid element, which represents a subdivision of the space, is a binary random variable. Either the corresponding space is occupied by some object or obstacle (hence it is attributed a value of 1), or it is empty (and has a value 0). The elements of the discretized ground truth grid (referred to as $O(\mathbf{x})$) are either 0 or 1. A grid estimate $\hat{O}(\mathbf{x})$, however, contains the probability that each grid element is occupied. Hence, each element can take any value between 0 and 1, while a value of 0.5 means that the state is truly uncertain. Fig. 5.1 shows an example of an occupancy grid map with coarse resolution for illustration.

The core of the binary Bayesian filter for recovering a grid map estimate from external observations is the conditional independence assumption between all grid elements o of O , and between all observations $\mathbf{z}_{0:t}$. This enables a recursive implementation of the filter for updating the state belief on each grid element: $b(o) = p(o = 1 | \mathbf{x}_{0:t}, \mathbf{z}_{0:t})$, that will be written $b(o) = p(o | \mathbf{x}_{0:t}, \mathbf{z}_{0:t})$ for notation simplicity. Furthermore, the conventional filter updates log-odds ratios:

$$l_t = \log \frac{p(o | \mathbf{x}_{0:t}, \mathbf{z}_{0:t})}{1 - p(o | \mathbf{x}_{0:t}, \mathbf{z}_{0:t})},$$

instead of the belief probability, as it leads to numerically stable updates using simple additions, as shown below. One may note that l_t can take any value between $-\infty$ and $+\infty$. Besides, the




Figure 5.1: An example of an occupancy grid with coarse resolution. It illustrates the typical effect of integrating a single range measurement to an obstacle into an initial uncertain map.

occupancy probability can easily be determined from the latter with the formula:

$$b(o) = 1 - \frac{1}{1 + \exp l_t}.$$

Here, we recall how the recursive filtering is derived using inverse measurement models $p(o|\mathbf{z})$, as presented in [94]. To do so, we start from the Bayes' formula:

$$\begin{aligned} p(o|\mathbf{x}_{0:t}, \mathbf{z}_{0:t}) &= \frac{p(\mathbf{z}_t|o, \mathbf{x}_{0:t}, \mathbf{z}_{0:t-1})p(o|\mathbf{x}_{0:t}, \mathbf{z}_{0:t-1})}{p(\mathbf{z}_t|\mathbf{x}_{0:t}, \mathbf{z}_{0:t-1})} \\ &= \frac{p(\mathbf{z}_t|o)p(o|\mathbf{x}_{0:t-1}, \mathbf{z}_{0:t-1})}{p(\mathbf{z}_t|\mathbf{x}_{0:t}, \mathbf{z}_{0:t-1})}, \end{aligned}$$

where the assumption of conditional independence between the observations is used. Now, after applying the Bayes rule to the measurement model:

$$p(\mathbf{z}_t|o) = \frac{p(o|\mathbf{z}_t)p(\mathbf{z}_t)}{p(o)},$$

we obtain:

$$p(o|\mathbf{x}_{0:t}, \mathbf{z}_{0:t}) = \frac{p(o|\mathbf{z}_t)p(\mathbf{z}_t)p(o|\mathbf{x}_{0:t-1}, \mathbf{z}_{0:t-1})}{p(o)p(\mathbf{z}_t|\mathbf{x}_{0:t}, \mathbf{z}_{0:t-1})}.$$

By applying the same reasoning for the probability of non-occupancy, one can obtain:

$$p(\neg o|\mathbf{x}_{0:t}, \mathbf{z}_{0:t}) = \frac{p(\neg o|\mathbf{z}_t)p(\mathbf{z}_t)p(\neg o|\mathbf{x}_{0:t-1}, \mathbf{z}_{0:t-1})}{p(\neg o)p(\mathbf{z}_t|\mathbf{x}_{0:t}, \mathbf{z}_{0:t-1})}.$$

Dividing the two previous equations leads to cancellation of probabilities that are difficult to calculate:

$$\begin{aligned} \frac{p(o|\mathbf{x}_{0:t}, \mathbf{z}_{0:t})}{p(\neg o|\mathbf{x}_{0:t}, \mathbf{z}_{0:t})} &= \frac{p(o|\mathbf{z}_t)p(o|\mathbf{x}_{0:t-1}, \mathbf{z}_{0:t-1})p(\neg o)}{p(\neg o|\mathbf{z}_t)p(\neg o|\mathbf{x}_{0:t-1}, \mathbf{z}_{0:t-1})p(o)}, \\ &= \frac{p(o|\mathbf{z}_t)}{1 - p(o|\mathbf{z}_t)} \frac{p(o|\mathbf{x}_{0:t-1}, \mathbf{z}_{0:t-1})}{1 - p(o|\mathbf{x}_{0:t-1}, \mathbf{z}_{0:t-1})} \frac{1 - p(o)}{p(o)}. \end{aligned}$$

After taking the logarithm of this expression, we obtain the recursive update formula of the Bayesian binary filter:

$$l_t = \log \frac{p(o|\mathbf{z}_t)}{1 - p(o|\mathbf{z}_t)} - \log \frac{p(o)}{1 - p(o)} + l_{t-1}.$$

5.2 Grid and feature-based mapping using ultrasonic guided waves

Our approach for mapping with combined representations builds on feature-based maps estimated via beamforming to construct inverse sensor models in an online fashion, so that a grid representation of a metal plate can be recovered within the OG framework. Simultaneously, the beamforming maps are filtered with binary masks to maintain consistency with the grid map, and at the same time, to mitigate the detrimental effect of interference.

5.2.1 Feature-based mapping via beamforming

As a reminder, let's briefly recall the feature-based mapping approach before highlighting some of its limitations. When a set of measurements have been acquired along a trajectory $\{x_i, y_i\}_{i=0..t}$, DAS beamforming is applied to recover the geometry of a metal plate from envelope signals e_i obtained through correlation of the measurement z_i with a wave propagation model \hat{g} . This yields 2D beamforming maps that are computed with:

$$\mathcal{L}_t(r, \alpha) = \sum_{i=1}^t e_i(d_i(r, \alpha)) = \sum_{i=0}^t e_i(|x_i \cdot \cos \alpha + y_i \cdot \sin \alpha - r|), \quad (5.1)$$

where (r, α) represent the polar coordinates of a candidate line boundary with equation: $x \cdot \cos \alpha + y \cdot \sin \alpha - r = 0$ expressed in a fixed frame with origin O , while $\mathcal{L}_t(r, \alpha)$ indicates the likelihood of the line existence. Next, a feature-based estimate of the map $\hat{\mathbf{M}}_t = \{r_l, \alpha_l\}_{l=1..4}$ is recovered by extracting the local maxima of \mathcal{L}_t through the resolution of the optimization problem:

$$\hat{\mathbf{M}}_t = \arg \max_{\mathbf{M}} \mathcal{L}_t(\mathbf{M}) = \arg \max_{\mathbf{M}} \sum_{l=1}^4 \mathcal{L}_t(r_l, \alpha_l),$$

where \mathbf{M} is restricted to be a rectangle. Note that an efficient resolution approach is proposed in the previous chapter by taking this constraint into account. Despite its efficiency in high signal-to-noise (SNR) conditions, the geometry retrieval can often be ambiguous due to interference, as an infinity of lines can equally account for one range measurement. Put in other words, assuming a high correlation $e_i(r_0)$, the application of DAS beamforming results in all lines (r, α) satisfying $|x_i \cdot \cos \alpha + y_i \cdot \sin \alpha - r| = r_0$ to have increased likelihood $\mathcal{L}_t(r, \alpha)$, whereas it is likely that the echo that explains the high value of $e_i(r_0)$ originates from reflection on a single line. In addition, the effect of interference is further exacerbated by the presence of multiple echoes and high-order reflections that are not considered in Eq. (5.1). Fig.5.2a illustrates the interference effect in a simulated scenario, where a mobile unit is driven to acquire acoustic measurements at predetermined positions marked as red points. As most of the measurements are acquired on a line parallel to an edge, there is an ambiguity on which side the real edge is. This results in the boundary on the right not being correctly identified from the beamforming map shown in Fig.5.2b.

The intuition behind the developed approach, which uses combined map representations, is that, when relying on detecting the earliest echo within the measurements, which indicates the distance to the closest edge, some information on the spatial environment can be directly

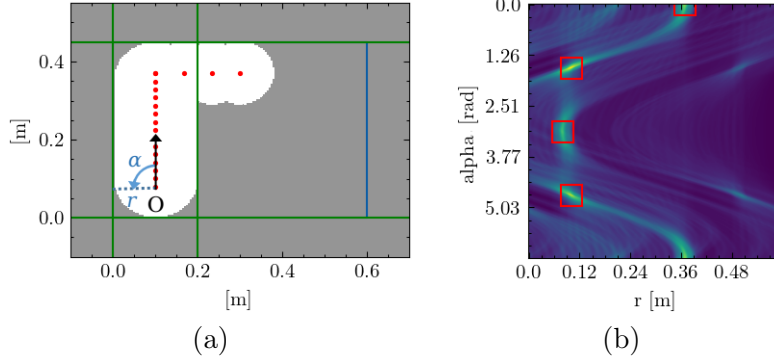


Figure 5.2: Illustration of the interference effect. (a) shows mapping results obtained in a simulated scenario, where one edge is not correctly identified from the beamforming map shown in (b). The origin O of the reference frame is indicated along with the (r, α) coordinate of the left line. The ground truth plate outline is represented by the blue rectangle.

integrated into the mapping results. First, it can be inferred that the disk centered at the sensor position whose radius equals the detected range contains only points on the current plate. Second, the lines that cross the total area identified as “on the plate” can be filtered out on the beamforming map to mitigate interference. This principle is illustrated in Fig. 5.3, where all the edges are now correctly identified after applying the mask filter to the beamforming map. The main issue of this approach is that the actual lines may be filtered as well when the measured ranges are overestimated.




Figure 5.3: Illustration of beamforming map filtering with the first acoustic echo for interference rejection. a) shows the mapping results obtained after applying, to the beamforming map in Fig. 5.2b, a binary mask that filters lines which can no longer exist, as they cross the area identified as “on the plate”. The on-plate area identified with the ground-truth echo measurements is represented in (a) as the lighter area. The origin O of the reference frame is indicated along with the (r, α) coordinate of the left line.

5.2.2 Estimation problem formulation

To obtain a more informative and time-consistent representation of the environment compared to a feature-based mapping only, our approach is designed to recover a grid representation of the 2D space,

$$O(\mathbf{x}) = \begin{cases} 0 & \text{if } \mathbf{x} = [x, y] \text{ is a point of the plate} \\ 1 & \text{otherwise} \end{cases},$$

by constructing an estimate \hat{O}_t with posteriors $\hat{O}_t(\mathbf{x}) = p(O(\mathbf{x}) = 1 | \mathbf{x}_{0:t}, z_{0:t})$ based on the detection of the closest edge. However, as the single-use of the detected range cannot be sufficient to recover areas outside the plate, the probability model shall rely on the feature-based map determined via beamforming to estimate the direction of arrival (DoA) of the detected echo. Besides, the introduced representation can be leveraged to filter out lines on $\mathcal{L}_t(r, \alpha)$ that are no longer relevant, because they cross the area identified as on the plate with \hat{O}_t , with the aim to mitigate interference. Hence, we formally introduce the binary mask filter associated with the ground-truth grid $O(\mathbf{x})$,

$$Q(r, \alpha) = \begin{cases} 0 & \text{if } \exists \mathbf{x} \text{ of the line } (r, \alpha) \text{ s.t. } O(\mathbf{x}) = 0 \\ 1 & \text{otherwise} \end{cases}, \quad (5.2)$$

so that the feature-based map $\hat{\mathbf{M}}_t$ can be retrieved from the masked beamforming map, where the mask is to be estimated as well: $\hat{\mathcal{L}}_t(r, \alpha) = \hat{Q}_t(r, \alpha)\mathcal{L}_t(r, \alpha)$. It is critical that such filtering is based on a probabilistic approach to avoid erroneously filtering lines. In practice, echo misdetection may occur due to noise and the propagation model's inaccuracy.

5.2.3 Binary filtering with adaptive inverse sensor models

The standard OG framework presented earlier is convenient for estimating grids of binary values. It relies on the essential assumptions that the measurements are conditionally independent given the grid state and that the states of the individual grid cells are conditionally independent given the measurements [94]. These assumptions, which are not always realistic, yield a computationally efficient algorithm that can run in real-time. Here, the filter involves an inverse measurement model that will be estimated from all past acoustic data based on the feature-based map retrieved via beamforming: $p(O(\mathbf{x}) | \mathbf{x}_t, e_t, \hat{\mathbf{M}}_t)$. More specifically, as $\hat{\mathbf{M}}_t$, which represents the parameters of the inverse model, is updated online, the interpretation of a measurement e_t can vary as well based on DoA estimates. In contrast, the standard OG framework would only involve $p(O(\mathbf{x}) | \mathbf{x}_t, e_t)$. Eventually, the grid can be recursively estimated with the standard Bayesian binary filter with log-odds ratios: $l_t = \log \frac{\hat{O}_t(\mathbf{x})}{1 - \hat{O}_t(\mathbf{x})}$ (and the assumption of a uniform prior):

$$l_t = l_{t-1} + \log \frac{p(O(\mathbf{x}) | \mathbf{x}_t, e_t, \hat{\mathbf{M}}_t)}{1 - p(O(\mathbf{x}) | \mathbf{x}_t, e_t, \hat{\mathbf{M}}_t)}, \quad (5.3)$$

with initial cell values $\forall \mathbf{x}, \hat{O}_0(\mathbf{x}) = 1/2$. For computational efficiency, the same approach is used to build \hat{Q}_t with a second inverse model $p(\hat{Q}_t(r, \alpha) | \mathbf{x}_t, e_t, \hat{\mathbf{M}}_t)$. Overall, the major challenge is specifying the two inverse models to determine the mapping strategy fully. In addition, this strategy shall be designed to maintain sufficient consistency between the grid \hat{O}_t and the mask \hat{Q}_t , as $\hat{Q}_t(r, \alpha)$ should formally depend on the full grid \hat{O}_t as in Eq. (5.2). However, in our approach, it is computed independently of the later given measurements and beamforming results.

One may observe that this approach is not perfectly rigorous mathematically speaking. Indeed, the use of Eq. (5.3) is based on the assumption of conditional independence between measurements (see how the binary filter formula is derived in the previous section 5.1), whereas $\hat{\mathbf{M}}_t$ is used as a conditional variable to reintroduce dependence with past observations. Yet, this mathematically inexact trick enables to leverage the simplicity of the binary filter with log-odd ratios for our problem. Another remark is that, here, we only consider the mapping problem, as the trajectory $\mathbf{x}_{0:t}$ is assumed known. Yet, our new mapping framework may seamlessly be integrated within a SLAM framework such as FastSLAM, as in Chapter 4.

5.2.4 Inverse model estimation in the Cartesian space

We first start the construction of a mixture model $p(O(\mathbf{x})|\mathbf{x}_t, e_t, \hat{\mathbf{M}}_t)$ through the modeling of imperfect echo measurements. Let ρ_t be the range to the closest edge retrieved from e_t . We assume that the echo detection can completely fail with a probability γ to construct a model robust to measurement error. In that event, the cell distribution shall be updated with a uniform distribution. With a probability $1 - \gamma$, the range is detected with a precision consistent with the wavelength. This is accounted for by introducing the model $p(\rho|e_t)$ that we model with a truncated Gaussian distribution $p(\rho|e_t) = \mathcal{N}_{[0, d_{\max}]}(\rho|\rho_t, \sigma^2)$ of support $[0, d_{\max}]$,

$$p(\rho|e_t) = \begin{cases} \frac{1}{\eta\sqrt{2\pi}\sigma^2} \exp\left\{-\frac{(\rho-\rho_t)^2}{2\sigma^2}\right\}, & \text{if } \rho \in [0, d_{\max}] \\ 0 & \text{otherwise} \end{cases},$$

where η is the normalizing factor and d_{\max} is the maximum detectable range that is introduced due to the finite observation window of the signals. The variable ρ denotes possible values for the range. Hence, the inverse model can first be expressed as a mixture and using a convolution:

$$p(O(\mathbf{x})|\mathbf{x}_t, e_t, \hat{\mathbf{M}}_t) = \frac{\gamma}{2} + (1 - \gamma) \int p(O(\mathbf{x})|\mathbf{x}_t, \hat{\mathbf{M}}_t, \rho) p(\rho|e_t) d\rho. \quad (5.4)$$

To build the model $p(O(\mathbf{x})|\mathbf{x}_t, \hat{\mathbf{M}}_t, \rho)$ which would state how the probability is updated for a given detected range ρ , we consider two scenarios. In the first one, the DoAs provided by $\hat{\mathbf{M}}_t$ would not be taken into account because the detected range would not be consistent with the expected range: *i.e* the map $\hat{\mathbf{M}}_t$ is assumed to be erroneous as it does not contain any line that is consistent with the range measurement, given that echo misdetection is already modeled with the probability γ . To account for this scenario, we introduce a model: $p_1(O(\mathbf{x})|\mathbf{x}_t, \rho) = h(|\mathbf{x}_t - \mathbf{x}| - \rho)/2$ where h denotes the Heaviside's function, $h = \mathbf{1}_{\rho > 0}$. Hence, only the points closer than the presumed range ρ can be marked as "on the plate", while those further away would have unchanged state belief distribution. Let us first assume that $p(O(\mathbf{x})|\mathbf{x}_t, \hat{\mathbf{M}}_t, \rho)$ is represented only by this model. As it is convolved in Eq. (5.4) with the truncated Gaussian distribution, the distribution updates are weighted with the range likelihoods. The convolution can seamlessly be expressed with the error function:

$$\int h(|\mathbf{x}_t - \mathbf{x}| - \rho) \mathcal{N}_{[0, d_{\max}]}(\rho|\rho_t, \sigma^2) d\rho = \frac{\Phi_t(|\mathbf{x}_t - \mathbf{x}|) - \Phi_t(0)}{\Phi_t(d_{\max}) - \Phi_t(0)}; \Phi_t(\rho) = \text{erf}\left(\frac{\rho - \rho_t}{\sigma\sqrt{2}}\right).$$

The mathematical derivation can be found in the appendix at the end of this thesis. With the current model, however, only points that are believed to be on the current plate are being recovered, and there is no mechanism for estimating areas outside the plate. To remedy that, we introduce a second model that will be used to integrate the information from $\hat{\mathbf{M}}_t$. Let's consider that the retrieval of a range ρ might originate from a reflection of the incident wave on an edge with polar coordinates $(\hat{r}, \hat{\alpha})$ from $\hat{\mathbf{M}}_t$. We introduce a second model to update the cells in the specific direction $\hat{\alpha}$ and beyond the range ρ ,

$$p_2(O(\mathbf{x})|\mathbf{x}_t, (\hat{r}, \hat{\alpha}), \rho) = \frac{1}{2} \beta(\hat{r}, \hat{\alpha}, \rho, \mathbf{x}_t) \times h((x - x_t) \cos \hat{\alpha} + (y - y_t) \sin \hat{\alpha} - \rho),$$

so that the cell "out-of-plate" probability can also increase. As the range-line association is not necessarily correct, we have introduced an association coefficient β between 0 and 1 that will make the probability update effective only when the considered range ρ matches with that expected

from a reflection on the edge $(\hat{r}, \hat{\alpha})$, and with Gaussian distribution to model uncertainty on the line parameters:

$$\beta(\hat{r}, \hat{\alpha}, \rho, \mathbf{x}_t) = \exp \left\{ -\frac{1}{2b^2} (|x_t \cdot \cos \hat{\alpha} + y_t \cdot \sin \hat{\alpha} - \hat{r}| - \rho)^2 \right\}.$$

For simplicity, the two aforementioned sub-models are then assembled together through averaging, yielding for a single line of the feature-based map:

$$p(O(\mathbf{x})|\mathbf{x}_t, (\hat{r}, \hat{\alpha}), \rho) = \frac{p_1(O(\mathbf{x})|\mathbf{x}_t, \rho) + p_2(O(\mathbf{x})|\mathbf{x}_t, (\hat{r}, \hat{\alpha}), \rho)}{2}.$$

This will result in Eq. (5.4) in the convolution of the model p_2 with the truncated Gaussian which can also be expressed with the error function:

$$\begin{aligned} \int \beta(\hat{r}, \hat{\alpha}, \rho, \mathbf{x}_t) h((x - x_t) \cos \hat{\alpha} + (y - y_t) \sin \hat{\alpha} - \rho) \times \mathcal{N}_{[0, d_{\max}]}(\rho | \rho_t, \sigma^2) d\rho \\ = \frac{\Gamma_{t, \hat{r}, \hat{\alpha}}((x - x_t) \cos \hat{\alpha} + (y - y_t) \sin \hat{\alpha}) - \Gamma_{t, \hat{r}, \hat{\alpha}}(0)}{\Phi(d_{\max}) - \Phi(0)} \end{aligned}$$

where we have:

$$\begin{aligned} \Gamma_{t, \hat{r}, \hat{\alpha}}(\rho) &= \frac{\lambda_t(\hat{r}, \hat{\alpha}) \cdot \sigma_{\text{new}}}{\sigma} \times \text{erf} \left(\frac{\rho - \mu_{\text{new}}}{\sigma_{\text{new}} \sqrt{2}} \right), \quad \mu_{\text{new}} = \frac{b^2 \rho_t + \sigma^2 d_t(\hat{r}, \hat{\alpha})}{b^2 + \sigma^2} \\ \sigma_{\text{new}} &= \frac{\sigma b}{\sqrt{b^2 + \sigma^2}} \text{ and } \lambda_t(\hat{r}, \hat{\alpha}) = \exp \left(-\frac{(\mu_{\text{new}} - \rho_t)^2}{2\sigma^2} - \frac{(\mu_{\text{new}} - d_t(\hat{r}, \hat{\alpha}))^2}{2b^2} \right). \end{aligned}$$

How this expression is derived can be found in the appendix at the end of the thesis.

Yet, we have so far considered that only one line is available, whereas the full feature-based map $\hat{\mathbf{M}}_t$ possesses four lines. Hence, for updating the state of the cell $\hat{O}_t(\mathbf{x})$, we take the maximum value of the convolution terms over the lines, so that the probability that \mathbf{x} is outside the plate can be increased as soon as there is at least one line estimate that is consistent with the range measurement. Furthermore, we introduce a multiplicative coefficient:

$$\mathcal{C}(\mathbf{x}, \mathbf{x}_t, \hat{r}_l, \hat{\alpha}_l) = \exp\{-\zeta|(y - y_t) \cdot \cos(\hat{\alpha}_l) - (x - x_t) \cdot \sin(\hat{\alpha}_l) - \hat{r}_l|\},$$

so that the update can be made effective locally around the axis passing through \mathbf{x} and its projection on the line $(\hat{r}_l, \hat{\alpha}_l)$. All brought together, the full inverse model in the Cartesian space can be expressed with:

$$\begin{aligned} p(O(\mathbf{x})|\mathbf{x}_t, e_t, \hat{\mathbf{M}}_t) &= \frac{\gamma}{2} + \frac{1 - \gamma}{4} \times \frac{\Phi_t(|\mathbf{x}_t - \mathbf{x}|) - \Phi_t(0)}{\Phi_t(d_{\max}) - \Phi_t(0)} \\ &+ \frac{1 - \alpha}{4} \max_{(\hat{r}_l, \hat{\alpha}_l) \in \hat{\mathbf{M}}_t} \left\{ \mathcal{C}(\mathbf{x}, \mathbf{x}_t, \hat{r}_l, \hat{\alpha}_l) \times \frac{\Gamma_{t, \hat{r}_l, \hat{\alpha}_l}((x - x_t) \cos \hat{\alpha}_l + (y - y_t) \sin \hat{\alpha}_l) - \Gamma_{t, \hat{r}_l, \hat{\alpha}_l}(0)}{\Phi_t(d_{\max}) - \Phi_t(0)} \right\}. \end{aligned}$$

The first term of the sum accounts for echo misdetection. The second one updates on-plate points with the detected echo up to a certain precision. More specifically, given a detected range, the area situated at a lower distance will have a decreased probability of being outside the plate. The last term recovers areas outside the plate based on the estimated feature-based map, which provides DoA estimates: for every line $(\hat{r}, \hat{\alpha})$ of $\hat{\mathbf{M}}$, the points further away from the sensors than the detected range, and following a direction towards the line, have increased probability of being




Figure 5.4: Illustration of the inverse model in the Cartesian space for different values of detected range (which is equal to the radius of the light disk surface). The sensor position is represented by the red point, and the considered line estimate is in green.

outside-the-plate when the range measurement matches the expected one. Besides, the update is made effective based the multiplicative coefficient $\mathcal{C}(\mathbf{x}, \mathbf{x}_t, \hat{r}, \hat{\alpha})$. It restricts the updated grid cells to limit mapping errors when the orientation of the line estimate is not stabilized.

To illustrate the effect of the inverse model, we show, in Fig. 5.4, the grid $p(O|\mathbf{x}_t, e_t, (\hat{r}, \hat{\alpha}))$ for different values of detected range ρ_t , and for a fixed and single line estimate $(\hat{r}, \hat{\alpha})$. In Fig. 5.4.a, the detected range value matches the distance to the line estimate. Thus, the probability to lie outside the plate, assigned to the points in the direction towards the line, is substantially increased based on the coefficient $\mathcal{C}(\mathbf{x}, \mathbf{x}_t, \hat{r}, \hat{\alpha})$. In Fig. 5.4.b, The probability increase is still noticeable, but it is weaker compared to the previous situation, as the match between the two distances is slightly worst. Last, when the difference between the measured and expected distances is too large as in Fig. 5.4.c, there are no cells that have increased probability of being outside the plate due to the lack of consistent DoA estimate.

5.2.5 Inverse model estimation in the line boundary space

As the mask \hat{Q}_t is determined independently of the grid \hat{O}_t with another binary filter, it is necessary to ensure sufficient consistency between the two grids, so that the filtered lines are consistent with the estimated on-plate area. We propose an inverse model in the line space which maintains exact agreement in range and approximate consistency in angle from the former inverse model:

$$p(Q(r, \alpha)|\mathbf{x}_t, e_t, \hat{\mathbf{M}}_t) = \frac{\gamma}{2} + \frac{1-\gamma}{4} \times \frac{\Phi_t(d_t(r, \alpha)) - \Phi_t(0)}{\Phi_t(d_{\max}) - \Phi_t(0)} + \frac{1-\gamma}{4} \max_{(\hat{r}_l, \hat{\alpha}_l) \in \hat{\mathbf{M}}_t} \left\{ \mathcal{C}'(\alpha, \hat{\alpha}_l) \times \frac{\Gamma_{t, \hat{r}_l, \hat{\alpha}_l}(|x_t \cdot \cos \hat{\alpha}_l + y_t \cdot \sin \hat{\alpha}_l - \hat{r}_l|) - \Gamma_{t, \hat{r}_l, \hat{\alpha}_l}(0)}{\Phi_t(d_{\max}) - \Phi_t(0)} \right\}$$

where $\mathcal{C}'(\alpha, \hat{\alpha}_l) = \exp\{-\mu|\alpha - \hat{\alpha}_l|\}$. The exact consistency in range is indeed maintained, as any line in the (r, α) space is uniquely mapped to the orthogonal projection of the robot position on the line at each measurement step. Such an approach is convenient, as the overall algorithm is computationally efficient, and can run in real-time. Finally, we decide to filter out, on the beamforming maps, only the lines whose nonexistence is certain enough. Thus, the features $\hat{\mathbf{M}}_t$ are retrieved from $\hat{\mathcal{L}}_t(r, \alpha) = \mathbf{1}_{\hat{Q}_t(r, \alpha) > \tau} \cdot \mathcal{L}_t(r, \alpha)$ where τ is chosen below 0.5.

5.2.6 Algorithm implementation

To summarize, we provide in Algorithm 3, the overall pipeline for implementation of our mapping approach with combined map representations.

Algorithm 3: Mapping_with_combined_representations($\mathbf{x}_t, e_t, \hat{O}_{t-1}, \hat{Q}_{t-1}, \mathcal{L}_{t-1}$)

Data: Position at current time \mathbf{x}_t , current envelope measurement e_t , grid map \hat{O}_{t-1} , grid mask filter \hat{Q}_{t-1} , beamforming map \mathcal{L}_{t-1} .

Result: Grid map \hat{O}_t , mask filter \hat{Q}_t and geometry estimate $\hat{\mathbf{M}}_t$ at current time t .

// Updates the beamforming map and retrieves a feature-based map estimate ;

$\mathcal{L}_t(r, \alpha) = \mathcal{L}_{t-1}(r, \alpha) + e_t(|x_t \cos \alpha + y_t \sin \alpha - r|), \forall (r, \alpha) ;$

$\hat{\mathcal{L}}_t(r, \alpha) = \mathbf{1}_{\hat{Q}_{t-1}(r, \alpha) > \tau} \cdot \mathcal{L}_t(r, \alpha) \forall (r, \alpha) ;$

$\hat{\mathbf{M}}_t = \arg \max_{\mathbf{M}} \sum_{(r_l, \alpha_l) \in \mathbf{M}} \hat{\mathcal{L}}_t(r_l, \alpha_l) ;$

// Updates the grid map and grid filter with the new inverse models ;

Determines ρ_t as the position of the first local maxima in e_t .

$l_t(\mathbf{x}) = \log \frac{\hat{O}_{t-1}(\mathbf{x})}{1 - \hat{O}_{t-1}(\mathbf{x})} + \log \frac{p(O(\mathbf{x})|\mathbf{x}_t, e_t, \hat{\mathbf{M}}_t)}{1 - p(O(\mathbf{x})|\mathbf{x}_t, e_t, \hat{\mathbf{M}}_t)} \forall \mathbf{x} ;$

$\tilde{l}_t(r, \alpha) = \log \frac{\hat{Q}_{t-1}(r, \alpha)}{1 - \hat{Q}_{t-1}(r, \alpha)} + \log \frac{p(Q(r, \alpha)|\mathbf{x}_t, e_t, \hat{\mathbf{M}}_t)}{1 - p(Q(r, \alpha)|\mathbf{x}_t, e_t, \hat{\mathbf{M}}_t)} \forall (r, \alpha) ;$

$\hat{O}_t(\mathbf{x}) = 1 - \frac{1}{1 + \exp l_t(\mathbf{x})} \forall \mathbf{x} ; \hat{Q}_t(r, \alpha) = 1 - \frac{1}{1 + \exp \tilde{l}_t(r, \alpha)} \forall (r, \alpha) ;$

return $\hat{O}_t, \hat{Q}_t, \hat{\mathbf{M}}_t$.

5.3 Numerical and experimental results

We assess the performance of our approach in different scenarios. In the first one, we use the ultrasonic measurements acquired on our aluminum plate of dimensions 600x450x6mm following the acquisition procedure presented in Chapter 3. As we only tackle the mapping problem here, only ground truth sensor positions will be used. In the second scenario, the data are obtained in simulation for plates of varying sizes and sensor trajectories. In the last scenario, we use the ultrasonic measurements acquired on our steel plate of dimensions 1700x1000x6mm, following the acquisition procedure presented in the previous chapter. For all the trials, the parameters of the filter are: $\gamma = 0.7$, $\sigma = 1\text{cm}$, $b = 2\text{cm}$, $\zeta = 3$, $\mu = 2$ and $\tau = 0.2$. The values of σ and b are consistent with the acoustic wavelength, approximately 2cm in our setups. The parameter γ is set with a high value to limit the effect of one measurement update and increase the algorithm's robustness w.r.t. echo misdetection. The detection of the range to the closest edge is achieved with a standard peak detector applied to the correlation signals. The resulting performance of our approach is compared with that of a more standard mapping strategy [62] as a baseline for comparison. The latter is slightly adjusted to better match our case-study: as in [62], all the points within a given range measurement ρ_t are directly marked as "on the plate". The lines, on the beamforming map, that cross the identified area are then filtered out, and a geometry estimate is then extracted from the processed beamforming result.

5.3.1 Results with hand-picked measurements

Fig. 5.5 shows the qualitative mapping results obtained with our approach at different steps of the scenario. The range to the closest edge is retrieved during the first step, but its orientation

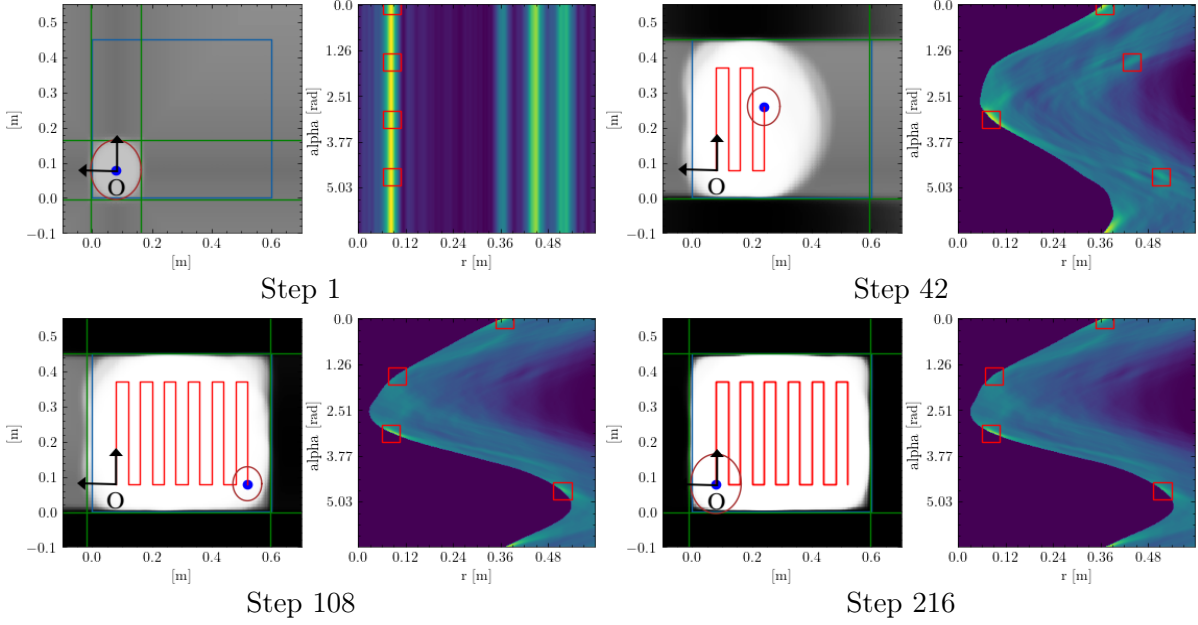


Figure 5.5: Mapping results achieved on the aluminum plate during Steps 1, 42, 108 and 216. Each sub-figure depicts, on the left, the sensor trajectory along with the detected range to the closest edge represented by the circles, and on the right, the filtered beamforming maps during the same step, where the red rectangles indicate the retrieved edge estimates.

is ambiguous, resulting in a squared plate estimate. During Step 42, an essential portion of the plate has been correctly identified as “on the plate”. Besides, the top and bottom sides of the plate have been marked as “outside the plate” based on the edge estimates retrieved from the filtered beamforming map. Nevertheless, the left edge has not been recovered because of interference that caused a line further away to be selected, as it cannot be filtered with our approach. During Step 108, all the edges are correctly retrieved from the filtered beamforming map. However, the area outside the plate on the left side has not yet been recovered because the edge estimation was incorrect when the mobile unit was close to it. This is achieved after simulating backwards the same trajectory, where the mapping is entirely achieved, as shown in Step 216. Ultimately, despite an average error of 4.3cm on range detection, the mean estimation errors on the line parameters are 0.5cm for the ranges r_l , and are inferior to one degree for the orientations α_l , $l = 1 \dots 4$. Furthermore, the two map representations can be considered coherent with each other, as only 7% of the plate points have been incorrectly identified as outside it.

Hence, the results demonstrate the efficiency of our approach for mapping a plate structure with combined representations, where the grid can account for plate edges that are recovered at any time of the mapping process. However, the detected edges often lie at the limit of the filtered area. To assess the robustness of our approach, we evaluate the final average mapping errors over 10 repetitions of the lawnmower trajectory for varying noise levels on the echo detection. This is achieved by adding Gaussian noise to the ground truth range, determined based on the known acquisition positions. We do the same experience by using the baseline approach. Fig. 5.6 shows the results. It can be noticed that, with our approach, a range estimation error close to 1 cm and an orientation error of a few degrees at most are maintained for noise levels comparable to the wavelength. This is not the case with the baseline approach, where the estimation errors increase rapidly with the noise level, which illustrates its lack of robustness to map the inside of the plate with error-prone measurements.

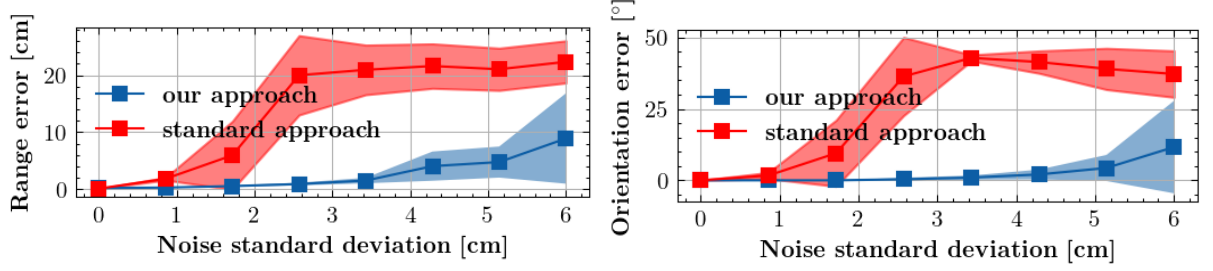


Figure 5.6: Final mapping errors w.r.t. the noise level on the echo detection assessed over 10 repetitions of the lawnmower trajectory. The errors are determined both for our approach and for the baseline one. The left plot shows the range errors, while the orientation errors are shown on the right plot. The solid lines represent the average values, while the upper and lower bounds of the colored areas are situated respectively above and below one standard deviation to the mean values.

5.3.2 Simulation results

To demonstrate that our method can generalize to various setups, we run it in simulation using data from three plates of different dimensions, and for a noise level of 5cm on echo detection. For comparison, we also repeat the process with the baseline approach, where all the points within the noisy range measurement are directly marked as "on the plate", and the lines, on the beamforming map, that cross the identified area are directly filtered out. We provide the resulting average errors and standard deviations assessed over 50 repetitions of random paths (at constant velocity) in Table 5.1. It can be observed that, for the three scenarios, our approach achieves a higher precision than the baseline approach.

Overall, the results demonstrate the effectiveness of our method for mapping a plate while filtering beamforming maps to reduce the effect of interference (without completely alleviating it), even in the presence of echo detection errors.

Plate size [mm]	Model	Range error [cm]	Angle error [degree]
600×450×6	Baseline	19.38 ± 4.83	33.46 ± 10.10
	Ours	0.17 ± 0.09	0.44 ± 0.37
1000×600×6	Baseline	31.73 ± 14.72	13.62 ± 15.10
	Ours	0.12 ± 0.11	0.37 ± 0.25
2000×1000×6	Baseline	69.62 ± 17.90	0.58 ± 0.89
	Ours	1.59 ± 3.59	0.32 ± 0.2

Table 5.1: Average estimation errors and standard deviations on the lines parameters obtained in simulation, using our approach (in **bold**) and the baseline method.

5.3.3 Results with a magnetic crawler

We evaluate our approach with a real robotic system and a large plate to demonstrate its viability for a practical robotic inspection task. In our experimental setup shown in Fig. 5.7a, the robot is equipped with an onboard accelerometer to provide accurate heading measurements, and with a single transducer for both acoustic emission and reception. Also, a pump is activated to continuously bring water through the tether at the interface between the sensor and the plate

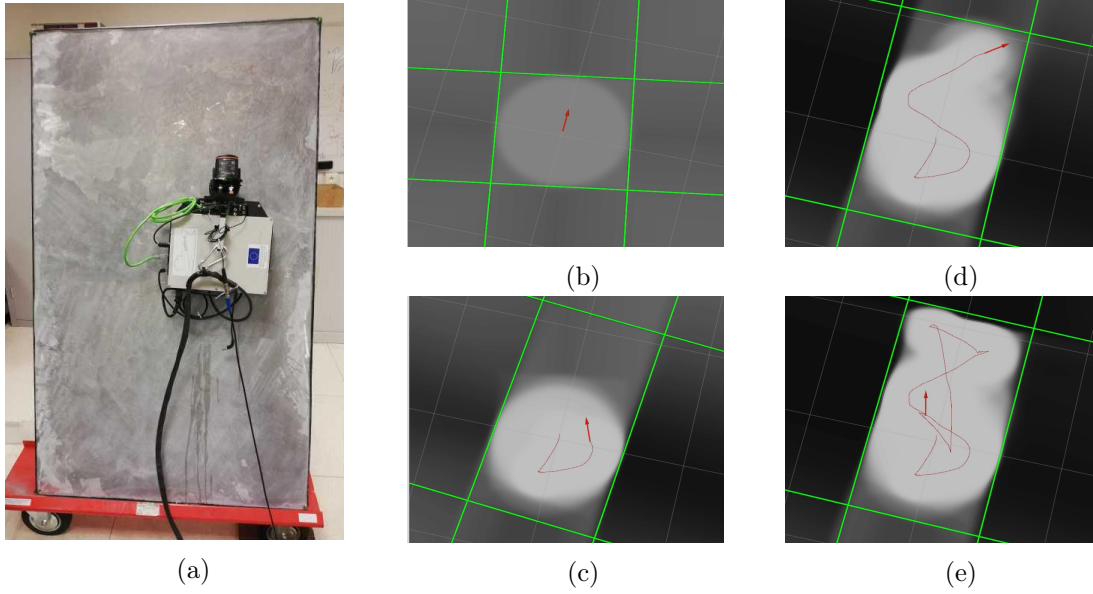


Figure 5.7: Combined feature-based and grid-based mapping with a robotic platform relying on UGWs. (a) shows the experimental setup. (b), (c), (d) and (e) present the mapping results at different steps along the robot transect. The robot pose (red arrow) and trajectory (red line) are also represented.

for sufficient coupling. During the acquisition, the robot is driven along a trajectory on the plate positioned nearly vertically, and the acquisition of one measurement is triggered each time the sensor has traveled 5 cm. For simplicity, we test our approach by assuming flawless odometry data, even though it is imperfect, and there may be slight errors in the kinematic model. However, the mapping procedure is expected to integrate seamlessly within a SLAM framework. Figures 5.7b, 5.7c, 5.7d and 5.7e show the qualitative mapping results obtained with our approach at different steps of the robot transect. It can be observed that, at the early stage, the size of the plate is almost recovered as the estimated dimensions are $1798 \times 1012mm$ during step (c), and the estimated orientation w.r.t. the robot is consistent with the reality. Besides, the inside and outside of the plate are correctly recovered, except for the regions near the corners, incorrectly marked as outside the plate. This slight inconsistency might be caused by the slight odometry errors and the increased complexity of the acoustic signals. Overall, these results demonstrate the applicability of our approach with a real robotic system. It may be useful for mapping a large structure made of several plates, as plate boundary detection is made robust, or by an active-sensing strategy for an optimal and autonomous mapping.

One of the limitations of our approach is that the parameters may be tuned carefully. Otherwise, proper lines may eventually be filtered on the beamforming map. In that situation, it is preferable to increase the noise parameters so that more measurements would be necessary before a line is filtered out. Also, our approach is evaluated on an isolated plate in a laboratory environment, while real metal structures present complex structural features such as welds or stiffeners. The constraint of rectangular shapes is not a limitation for mapping storage tanks or ship hulls, as they are almost entirely made of rectangular panels. However, extensions of our approach are to be investigated to work with more general structures. It is expected it would only require an extension of the optimization method presented in Chapter 4 to retrieve the

edges from the beamforming maps. At the same time, the design of the Bayesian framework could be left unchanged.

5.4 Towards ultrasound-based autonomous exploration

The algorithms presented so far were dedicated to localization and mapping purposes: following a *predetermined* sensor trajectory for collecting the ultrasonic measurements, the localization and mapping performance of the different methods was evaluated. Yet, if one seeks to maximize the reconstruction performance, the question of *how* the robot should move is also central. For example, some robot trajectories may result in better map reconstruction than others if the measurements positions are such that the retrieved quantity of information is maximal. In our case study, automatic path planning and path following can be beneficial to the development of a fully automated inspection agent that can meet reconstruction optimality criteria. One could aim at maximization of the reconstruction quality of a large structure for reliable identification of potential defects, or minimize the reconstruction time to reduce inspection costs.

As a first step towards such an application, we introduce, in this section, an active-sensing strategy based on frontier exploration [102] to enable the autonomous reconstruction of the geometry of a metal surface by a mobile robot relying on ultrasonic echoes. The method builds on our framework for mapping with combined representations: during the mapping process, the grid and feature-based map estimates of the unknown metal surface are leveraged to estimate the information gain provided by measurements made at potential next positions. Following a utility function that balances information gain and travel cost, the sensor is automatically directed towards the best frontier point that lies in the boundary between "explored" and "unexplored" areas. The developed method is evaluated in simulation, and is compared with multiple exploration strategies, essentially closest and random frontier point selection. Finally, an experiment using a mobile robot equipped with a co-localized pair of transducers is used to validate the viability of the approach. It is acknowledged that what follows is a mere summary of the method presented in [77].

5.4.1 Related works

Autonomous robotic exploration to map unknown environments is a major research field in robotics, aiming at algorithmic path planning and path following solutions to maximize information gain and minimize costs. Frontier-based approaches yield common solutions to the exploration problem, and are extensively used to map indoor environments [102, 27]. By moving towards new frontiers between mapped and unknown areas, a mobile robot can extend its map into new territory until the entire environment has been explored. The key to effective frontier exploration is the selection of target frontier points. In the original method [102], the closest frontier point is selected as the next target. In most cases, exploration strategies select the next best frontier by evaluating candidate locations according to different criteria, resulting in various extensions of the basic frontier-based exploration strategy [45, 63, 27]. In [98], for example, a utility function balancing the travel cost with the information gain is used.

There are currently different methods to evaluate the information gain: one is to directly assess the undetected space size in the nearby region of the target frontier point [4]. Another one is to calculate the information entropy to assess the uncertainty on map prediction [10, 94]. Yet, to the best of our knowledge, the use of such approaches for mapping using omnidirectional UGWs in a pulse-echo setup has not been thoroughly studied in the literature.

A Lamb wave-based exploration strategy for mapping a metal plate is proposed in [62] using a pitch catch setup. Yet, the transducers are manually positioned in the desired (and sometimes distant) positions, which does not leverage constant signal acquisition over the continuous robot displacement. Furthermore, the mapping method only relies on detection of the first reflection on the plate boundaries, and it is subject to gridlock situations which are not predictable.

5.4.2 Overview of the frontier exploration approach

Our approach relies on the mapping framework with combined representations presented earlier. During Step t of the mapping process, the grid map estimate \hat{O}_t is used to determine the K_t frontier points $\mathbf{p}_{t,k}, k = 1..K_t$, that lie in the frontier between explored and unexplored areas on the plate. Along with the feature-based map estimate $\hat{\mathbf{M}}_t$ obtained via DAS beamforming, the grid \hat{O}_t is used to evaluate the cumulative entropy:

$$I(\mathbf{p}_{t,k}, \hat{\mathbf{M}}_t, \hat{O}_t) = \sum_{\mathbf{p} \in V(\mathbf{p}_{t,k}, \hat{\mathbf{M}}_t)} E(\mathbf{p}),$$

of all the points in the vicinity $V(\mathbf{p}_{t,k}, \hat{\mathbf{M}}_t)$ of $\mathbf{p}_{t,k}$, that are expected to lie on the plate. $e(\mathbf{p})$ is the well known entropy over the state probability of a point \mathbf{p} of the grid:

$$E(\mathbf{p}) = -\hat{O}_t(\mathbf{p}) \log \hat{O}_t(\mathbf{p}) - (1 - \hat{O}_t(\mathbf{p})) \log(1 - \hat{O}_t(\mathbf{p})).$$

The overall utility of a frontier point is defined as a weighted sum between the resulting entropy value and the distance between the sensor and the frontier point:

$$U(\mathbf{p}_{t,k}, \hat{\mathbf{M}}_t, \hat{O}_t) = I(\mathbf{p}_{t,k}, \hat{\mathbf{M}}_t, \hat{O}_t) - \lambda \cdot \|\mathbf{p}_{t,k} - \mathbf{x}_t\|,$$

where λ is a positive parameter, and \mathbf{x}_t the robot position at time t . Hence, by selecting the frontier point with maximum utility as next best point to visit, a moving robot will be driven towards nearby locations where the reconstruction uncertainty is still important.

During the whole mapping process, the feature-based and grid-based estimation procedures are being executed. Thus, as the maps, frontier points and their respective utility vary accordingly, the selection of the best frontier point to move to is repeated at every measurement step. Finally, a frontier point is always attributed a utility of 0 when it is considered closer to a boundary than a fixed distance, so that other frontiers can also be explored. The frontier exploration is considered complete when there are no more frontier points to evaluate regarding their distance to the edge. An illustration of a mapping output provided by the proposed approach is presented in Fig. 5.8.

To summarize, the different steps of one exploration step are the following:

- Frontier points are extracted from the grid map.
- The utility of each frontier point is evaluated.
- The candidate pose with the highest utility is selected as the next goal.
- The robot navigates to the target position.




Figure 5.8: Illustration of a mapping output provided by our approach. The blue dot represents the sensor position and the blue line is its overall trajectory. The red rectangle represents the estimate $\hat{\mathbf{M}}_t$ obtained via DAS beamforming. The grid $\hat{\mathbf{O}}_t$ is accounted for with the color of the cells. The valid (resp. invalid) frontier points corresponds to the red (resp. black) dots. The green point indicates the selected next best point.

5.4.3 Overview of the exploration results

To assess the performance of our approach, we perform several tests conducted in simulation and with a real robot. During each experiment, as the mapping time increases, we keep track of the distance traveled by the robot, along with the ratio of covered surface of the plate as a performance criterion. The latter is defined as the ratio between the plate surface correctly labeled as "on the plate" and the entire plate surface. Contrary to the estimation errors on the plate geometry parameters, it is a good indicator of how spread the measurements positions are. Furthermore, it is relevant to aim at complete measurement coverage during a robotic inspection operation, as the entire surface of the structure is to be inspected. For both the simulated and real-world experiments, we will consider a flat rectangular metallic plate with dimensions 1700x1000x6mm.

First, we run our approach in a simulated scenario. In Fig. 5.9, we show the mapping outputs at different steps, along with the sensor trajectory resulting from the application of our exploration strategy. One can observe that, initially, the map is very uncertain. Thus, the selection of frontier point is essentially random. During Step 41, after a seemingly erratic trajectory due to high uncertainty on the map, three edges are correctly recovered. Yet, the mapping error on the grid is relatively important due to the right edge that is incorrectly recovered. Following the selected frontier points, the robot is driven to the right part of the plate, until all the points have been marked as invalid. This results in the correction of the errors in the grid map. During the last step, the sensor has followed the remaining frontier points until there is no valid point left, leading to termination of the exploration process with a seemingly low mapping error.

For quantitative comparison with baselines, we now also evaluate the coverage performance of three other exploration strategies: the classic frontier approach where the closest frontier point is systematically selected as next goal location, a random approach where the next frontier point is sampled from a uniform distribution, and our method using the true plate's geometry, which is expected to lead to the best coverage performance.

Each exploration algorithm is run for the same simulation setup: the metal plate has dimensions 1700x1000x6mm, and the sensor is initially placed at 20 cm from two boundaries, in a corner. To assess the repeatability of each exploration strategy, we evaluate the means along with the 10% and 90% quantiles of the mapping coverage. For better visualization, we only present coverage results up to 83% of mean coverage value. The results of the experiments can be seen in Fig. 5.10. It can be observed that the proposed approach outperforms the random

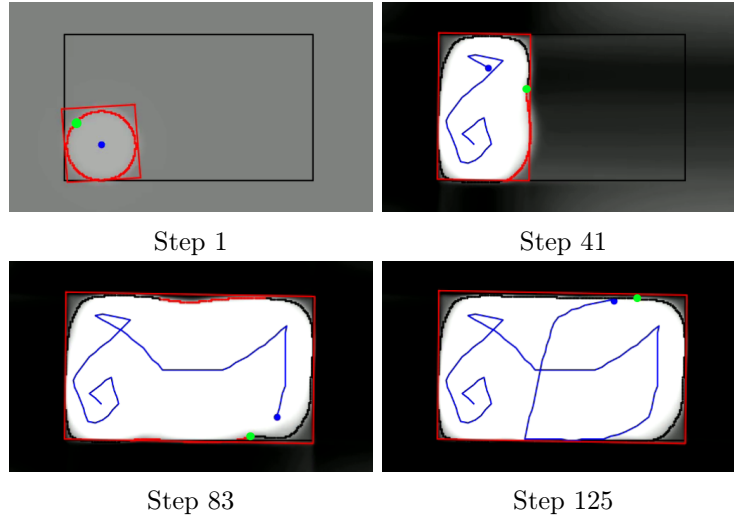


Figure 5.9: Mapping outputs obtained in simulation during different steps of the mapping process. The blue dot represents the sensor position, and the blue line is the sensor trajectory. The red rectangle represents the feature-based map estimate obtained via DAS beamforming. The grid \hat{O}_t is accounted for with the color of the cells. The valid (resp. invalid) frontier points correspond to the red (resp. black) dots. The green point indicates the selected next best point.

approach (which is the least efficient) and the standard closest point strategy, as higher mapping coverage values are obtained for lower traveled distances. The reason is that it takes both into account information gain and travel cost. Naturally, knowing the ground truth plate geometry leads to better results. It represents an upper bound on our approach performance.

Finally, the performance of our approach is experimentally validated using a real robotic platform. We use a turtlebot equipped with a structure for holding two transducers and maintaining them in contact with the surface, as depicted in Fig. 5.11. A layer of water is placed on the plate surface as a coupling medium. For signal acquisition, we use two-tone bursts of a sinusoidal wave at 100kHz as the excitation and a sampling frequency of 1,25 MHz for acquisition.

During all the experiments, we use a camera which provides a top view recording of each experiment. The tags depicted in the experimental setup shown in Fig. 5.11 are used to track the robot position on the plate using the AR tags tracking package available on ROS². This system is used here for ground truth positioning. Simultaneous localization, mapping and exploration is left for future work. Also, to evaluate our exploration system independently to the quality of the controller, the movement between waypoints is implemented with the joystick. Yet, the destination points are automatically generated by our algorithm.

In the next, we compare the performance of the closest frontier point selection and random exploration strategies with the proposed approach. We perform five repetitions of each algorithm, each time placing the sensors at the same starting position corresponding to $[x_0, y_0] = [25cm, 45cm]$ expressed in the plate frame. The results are provided in Fig. 5.12, where we represent the evolution of plate coverage up to a value of 91% for the mean coverage. One may observe that the coverage percentage increases relatively fast for the proposed method, as an average coverage value of 83% is reached with a displacement of 2.38m. We remark that the random approach outperforms our method at the beginning, but plateaus around 80% coverage due to its inability to select informative waypoints to further improve the mapping results.

²http://wiki.ros.org/ar_track_alvar




Figure 5.10: Evolution of the plate surface coverage for the four exploration strategies with respect to the traveled distance in the simulations. The discontinuous lines represent the coverage mean value over 50 repetitions of each method, while the upper and lower bounds of the colored areas represent respectively the 90% and 10% quantiles.




Figure 5.11: Top view of the experimental setup. A turtlebot platform is used along with AR tag markers for ground truth positioning.




Figure 5.12: Evolution of the plate surface coverage for the three exploration strategies with respect to the traveled distance during the experiments. The solid lines represent the coverage mean value over 5 repetitions of each method, while the upper and lower bounds of the colored areas represent respectively the 90% and 10% quantiles.

The closest point selection method is outperformed by the two others due to inefficiency of the greedy strategy. Overall, the results illustrate the potential of an active-sensing mapping system for automated inspection. Future works shall be directed towards implementation and testing in more realistic scenarios to further make possible the emergence of such an application.

5.5 Polygonal shapes reconstruction from acoustic echoes

In the previous sections, the mapping problem was limited to recovering the geometry of rectangular panels using UGWs. By leveraging DAS beamforming and the rectangular constraint, one can efficiently recover the position of the plate boundaries, as the explicit detection of all the echoes within the signals and their association with potential line boundaries is not required. As seen in Chapter 4, this allows for integration of the approach within a FastSLAM framework that can seamlessly run online. Yet, it is restricted to rectangular geometries.

This assumption may not be restrictive if one is only interested in recovering the geometry of surfaces made of metal panels, and to achieve SLAM on structures like ship hulls or storage tanks. However, reducing the number of assumptions on the map may be beneficial to a potential real-world inspection application. More specifically, defects, which may also act as acoustic scatterers, could be integrated in the map. Thus, they could be detected using the same DAS beamforming approach. In this situation, however, the rectangular constraint can no longer hold for recovering them. In addition, DAS beamforming could also be used for mapping more general structures, with more arbitrary geometries.

The objective of this last section is to propose a more general approach based on beamforming along with combined map representations to recover arbitrary 2-dimensional polygonal geometries from acoustic echoes. It applies DAS beamforming in the geometrical boundary space, and relies on the *single* detection of the closest boundary to map the inner surface of the geometric shape, using a grid representation. The feature-based geometry estimate is then constructed via minimization of a regularized cost function using a simulated annealing optimizer to select boundaries, on the beamforming results, consistent with the mapped inner space. We also de-

sign a boundary rejection criterion to recover their exact number based only on a specified upper bound. We test our approach in simulation for various geometries and actual conditions using experimental ultrasonic measurements acquired on a metal plate sample. The results show that it can successfully recover the geometric shape of the reverberant environment, the exact number of boundaries, and has the same performance as the baseline method, presented earlier, whose use is limited to recovering rectangular shapes.

5.5.1 Related works

The use of ultrasonic echoes for recovering polygonal shapes such as a room geometry, and using a single mobile unit, has been studied in various works that have been mentioned earlier [51, 53, 69]. However, these are only tested in simulation. They rely on essential assumptions, such as first-order echoes are directly available, the echo-wall association is being flawlessly known, or the number of walls to recover is being known a priori. Due to noisy measurements and inaccurate signal models, these assumptions may not always hold for a real-world scenario, resulting in low accuracy reconstruction, as illustrated by the experiments in Chapter 4.

When deployed in real-world scenarios, other acoustic methods for mapping indoor environments rely on directivity of the signal acquisition [79]. Another setup, which is more computationally involving, is to use an array of sensors for DoA estimation [80, 61]. Yet, the mentioned works only rely on the earliest signal in the measurements, whereas late reflections also convey information on the environment. Overall, the efficient mapping of arbitrary polygonal shapes from omnidirectional acoustic echoes is still to be established.

5.5.2 Problem statement

Let's first re-introduce the considered problem. We consider a setup where a single (collocated) emitter/receiver acoustic device is driven to acquire acoustic measurements $z_{0:T}$ at positions $\mathbf{x}_{0:T}$ on a bounded metal surface, where $\mathbf{x}_i = [x_i, y_i, \theta_i]^T$ contains the two-dimensional position and orientation of the robot. At the i -th scanning position, a pulse $s(t)$ is used to emit an omnidirectional ultrasonic wave. The received signal $z_i(t)$ contains the reflections of the incident wave on the plate boundaries. We also assume that, at each position \mathbf{x}_i , an accurate measurement (or estimate) of the range ρ_i to the closest boundary is available by processing the acoustic data. This measurement may only be subject to little noise in practice, as it may be one of the most energetic echoes. Based on all range measurements $\rho_{0:T}$, the surface A_T is obtained, as a grid map, using the unions of all disks centered at sensor positions \mathbf{x}_i , and whose radius equal to the detected range ρ_i . By construction, this area cannot be crossed by a boundary. For simplicity of the presentation, we assume that noise-free odometry information (*i.e.* linear and rotational displacement between the positions of the successive measurements) is available. Yet, the proposed method may be integrated in a FasSLAM algorithm as in Chapter 4 for localization.

Our objective is to recover the shape \mathbf{M} of the plate that we now assume polygonal, instead of the much more limiting assumption of a rectangular geometry that has been made so far, using ultrasonic measurements. One may note that not only we only address the mapping problem here, but we assume little noise on the measurement ρ_i . To alleviate the need of the latter assumption, one may integrate the proposed method within the mapping framework with combined representations presented in the first part of this chapter. Indeed, it has been observed that the mapping results were robust w.r.t echo detection errors.




Figure 5.13: Schematic of the problem setup and notations. An example of shape to estimate is represented by the red polygon. A candidate geometry estimate is represented by the set of green lines, whose (r, α) coordinates are only provided for one of them. The area A formed with a single range measurement ρ is represented by the lighter area. The inner surfaces of A and $\hat{\mathbf{M}}$ are delimited with the discontinuous dashed lines.

5.5.3 A regularized cost for polygonal shape reconstruction

We propose to estimate \mathbf{M} using the (now very classic) DAS beamforming approach, and only adapt the cost function whose minimization provides a feature-based geometry estimate. This is achieved by adding a regularization term to the initial cost function for maintaining consistency between the feature-based map estimate and a grid representation of the plate inner surface. The latter is obtained using the range measurements ρ_i .

The geometry \mathbf{M} to recover is assumed to be a 2D convex polygon that can be represented by its linear boundaries: $\mathbf{M} = \{r_l, \alpha_l\}_{l=1..L}$ where the parameters (r_l, α_l) define the line equation in the 2D plane, expressed in a fixed frame whose origin will be chosen as the initial robot position, and the x axis will be aligned with the initial robot heading. The number L of boundaries is assumed a priori unknown. The adopted notations are precised in Fig. 5.13.

First, the envelope measurements e_i are classically obtained by correlating the measurements z_i with a wave propagation model. Next, the beamforming map to assess the likelihood of existence for any line boundary (r, α) is determined with:

$$\mathcal{L}_T(r, \alpha) = \sum_{i=0}^T e_i (|x_i \cdot \cos \alpha + y_i \cdot \sin \alpha - r|),$$

Based on the former mapping approach, only rectangular shapes could be recovered by retrieving exactly four local maxima of \mathcal{L}_T through the maximization of their cumulative likelihood. We propose to extend this approach to recover, in principle, arbitrary 2D polygonal shapes. The estimated geometry has the form: $\hat{\mathbf{M}} = \{\hat{r}_k, \hat{\alpha}_k\}_{k=1..K}$, where the maximum number of boundaries K is such that $K \geq L$. This is a very minimal hypothesis, as K can be set much higher than the unknown L .

Next, for a given robot trajectory $\mathbf{x}_{0:T}$, we determine the inner surface A_T using the range measurements $\rho_{0:T}$. Based on the beamforming map \mathcal{L}_T that we normalize by its maximum value, we introduce the cost function:

$$J(\hat{\mathbf{M}}) = - \left(\sum_{(r_k, \alpha_k) \in \hat{\mathbf{M}}} \mathcal{L}_T(r_k, \alpha_k) + \lambda \frac{\hat{\mathbf{M}} \cap A_T}{\hat{\mathbf{M}} \cup A_T} \right), \quad (5.5)$$

where the first term aggregates the likelihood of presence of the estimated boundaries, and the second term is the well known Intersection over Union (IoU), widely used for computer vision purposes [75]. In Eq. (5.5), the surfaces $\hat{\mathbf{M}}$ and A_T are as described in Fig.5.13. The IoU term (that we will simply write as $\text{IoU}(\hat{\mathbf{M}}, A_T)$) is always between 0 and 1. It is zero when $\hat{\mathbf{M}}$ and A_T have an empty intersection, and it is 1 when the two surfaces are equal. λ is a positive parameter that balances the importance between the boundaries likelihood and IoU terms. Our geometry estimate is then defined as:

$$\hat{\mathbf{M}}_T = \arg \min_{\hat{\mathbf{M}}} J(\hat{\mathbf{M}}).$$

Hence, the minimization of the cost criterion leads to maximization of the cumulative likelihood of the estimated boundaries and maintains consistency between the estimated inner surface of $\hat{\mathbf{M}}$ with that of A_T . In other words, the IoU in Eq. (5.5) acts as a regularization term for selecting likely boundaries on the beamforming results.

5.5.4 Shape reconstruction with simulated annealing

An optimization strategy is required to recover a geometry estimate by minimizing the cost. Simulated annealing (SA) [49] is a relatively simple and efficient approach for estimating a global minimum of a non-linear cost function. This metaheuristic is inspired by metallurgy, where, to form a perfect crystal (which corresponds to the state of minimal energy), a pure liquid substance is slowly cooled. A random perturbation is applied to the current state at each iteration of the optimization process. The perturbation is systematically accepted if the energy is decreased. It is accepted according to a Boltzmann probability distribution when the energy is increased to escape local minima. The temperature, which is a hyperparameter that controls the acceptance rate, is decreased slightly after each iteration, so that the probability of increasing the energy goes to zero. Even though simulated annealing offers no guarantees that the optimum will be found, this strategy can lead to satisfying solutions in a relatively restricted number of iterations. In the following, we detail the different steps of our custom implementation for mapping arbitrary polygonal geometries from acoustic echoes.

Initialization

First, the geometry estimate is randomly initialized with $\hat{\mathbf{M}}_T = \{\hat{r}_k, \hat{\alpha}_k\}_{k=1..K}$, where the \hat{r}_k and $\hat{\alpha}_k$ are sampled from independent uniform distributions.

State disturbance

At iteration p of the optimization process, each boundary of the estimate is iteratively disturbed with a random process. For the k -th boundary, we sample:

$$\begin{cases} \tilde{r}_k = \max(\hat{r}_k + \Delta_r \cdot \chi_r^3, 0) \\ \tilde{\alpha}_k = \hat{\alpha}_k + \Delta_\alpha \cdot \chi_\alpha^3 \bmod 2\pi \end{cases},$$

where Δ_r and Δ_α are positive parameters, and χ_r and χ_α are independently sampled from uniform distributions: $\mathcal{U}([-1, 1])$. They are raised to the power 3 so that small variations are more likely. Still, large variations (*i.e* in the order of Δ_r and Δ_α) can be possible to explore regions of the search space that are far occasionally from the current state. Furthermore, to speed up the optimizer convergence, we determine:

$$(\tilde{r}_k, \tilde{\alpha}_k) \leftarrow \arg \max_{(r, \alpha) \in V(\tilde{r}_k, \tilde{\alpha}_k)} \mathcal{L}_T(r, \alpha),$$

where $V(\tilde{r}_k, \tilde{\alpha}_k)$ denotes the set of boundaries that lie in the neighborhood of $(\tilde{r}_k, \tilde{\alpha}_k)$ that is to be appropriately defined. The purpose of this step is to only select $(\tilde{r}_k, \tilde{\alpha}_k)$ that are local maxima, which is a property that is expected for the true boundaries. The former edge $(\hat{r}_k, \hat{\alpha}_k)$ in $\hat{\mathbf{M}}_T$ is then substituted to the sampled one, yielding a new geometry $\tilde{\mathbf{M}}_T$.

Simulated annealing acceptance process

The new candidate map is kept as the current state following an SA acceptance process. First, we determine the cost variation $\Delta J = J(\tilde{\mathbf{M}}_T) - J(\hat{\mathbf{M}}_T)$. Let's χ be a sample from $\mathcal{U}([0, 1])$. The candidate geometry is kept only if we have:

$$\chi < \exp\left(-\frac{\Delta J}{T(p)}\right),$$

where $T(p)$ is the temperature at iteration p of the optimizer. Hence, if the cost is decreased, the new map is systematically kept as the new current state $\tilde{\mathbf{M}}_T$. Otherwise, it is kept according to the Boltzmann distribution to escape local minima.

Edge filtering

As it is expected that the edges number parameter K is overestimated, some of them may not contribute significantly to the definition of the recovered geometry once the optimizer has converged. To derive a proper geometry with a correct number of boundaries, we apply a rejection criterion to remove non-informative edges. We iterate over each boundary $(\hat{r}_k, \hat{\alpha}_k)$ in $\hat{\mathbf{M}}_T$, and we remove it to yield a new candidate geometry $\tilde{\mathbf{M}}_T$. We then assess whether:

$$\left| \frac{\text{IoU}(\tilde{\mathbf{M}}_T, A_T) - \text{IoU}(\hat{\mathbf{M}}_T, A_T)}{\text{IoU}(\hat{\mathbf{M}}_T, A_T)} \right| < \epsilon,$$

where we will take $\epsilon = 10^{-2}$. If the IoU variation is lower than the threshold, then the k -th boundary can be removed, as its contribution to the definition of the geometric surface is considered negligible. Then, the geometry $\tilde{\mathbf{M}}_T$ can be taken as the new candidate $\hat{\mathbf{M}}_T$, and the process is iterated on the next boundary. Also, the rejection criterion is not applied if the candidate map $\hat{\mathbf{M}}_T$ has 3 edges (which is a lower bound for a polygonal shape).

5.5.5 Simulation results

We perform different experiments to assess the efficiency of our method for recovering arbitrary 2-dimensional polygonal geometries from acoustic echoes acquired over a mobile unit trajectory. We first consider three simulated scenarios, where the geometry to be recovered has either 3, 4 or 5 boundaries. The robot paths are randomly generated with successive linear and rotational displacements. The ultrasonic measurements acquired over the path are simulated using the Image source model, and the ultrasonic model g introduced in Chapter 2. It should be noted that such a data model leads to over-energetic high-order reflections, as acoustic absorption from the boundaries is not taken into account. The excitation $s(t)$ is a two tone bursts sinusoidal wave at 100 kHz. We add Gaussian noise on the measurements to maintain a fixed signal-to-noise ratio $\text{SNR} = 0$. The considered geometries and paths are provided in Fig. 5.14.

We run 10 repetitions of the optimizer for each scenario to assess the repeatability of our approach. The optimizer is set with parameters $\lambda = 3$, $T(p) = 5/p^{0.8}$, and $V(r, \alpha)$ is chosen as

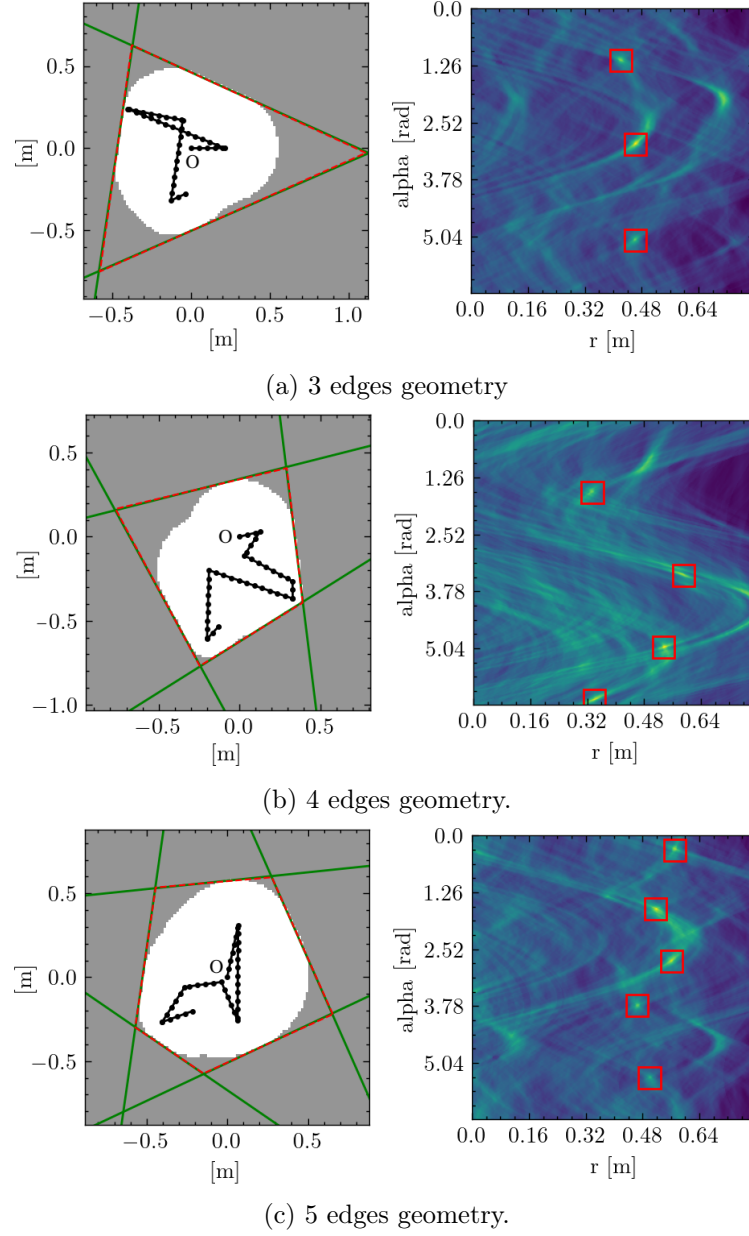


Figure 5.14: The different simulated setups along with an instance of mapping results. On the left of each sub-figure, the outline of the ground-truth geometry is provided in discontinuous red, the green lines account for the estimated edges, and the robot path and acquisition positions are represented in black. The corresponding beamforming maps are provided on the right, where the red rectangles indicate the retrieved edges. The edge coordinates are expressed in the fixed frame, whose origin is the initial sensor position O , and the x -axis is aligned with the initial robot heading.




Figure 5.15: Evolution of the performance criteria for the 3 scenarios, and for 10 repetitions of the optimizer. The top plots show, in blue, the evolution of the costs. The middle plots depict the estimated number of edges after application of the edge rejection criterion. The bottom plots show the IoU between estimated and ground-truth surfaces. The median values are in red.


the set of edges with line coordinates within the rectangular area centered in (r, α) , and with size $5.4\text{cm} \times 48^\circ$. Also, we use $K = 10$ as the maximum number of edges. Fig. 5.15 shows the evolution of the cost, as well as the number of estimated edges after use of the edge rejection criterion during 300 optimizer iterations, and for the 10 repetitions. We also display the evolution of $\text{IoU}(\hat{\mathbf{M}}, \mathbf{M})$, which is the IoU between the estimated inner surface (based on the estimated edges) and the ground truth surface \mathbf{M} . Hence, an IoU of 1 translates into a perfect geometry estimation. The results demonstrate that, for all the considered scenarios, our approach can seamlessly recover, after convergence, both the actual number of boundaries and the geometric shape of the simulated plates.

5.5.6 Experimental results


To validate our approach, we run it using the experimental ultrasonic data acquired on our $600 \times 450 \times 6\text{mm}$ rectangular aluminium metal plate by moving, by hand, two nearly collocated contact piezoelectric transducers for signal emission and reception. The optimizer is set with the *exact* same parameters than those used for the simulation tests to highlight the robustness of our approach. Also, we do not use ground-truth ranges $\rho_{0:T}$, but apply a standard peak detection method on the correlation signals to estimate them. This only leads to a slight estimation error due to high SNR. The considered simulated robot trajectory is depicted in Fig. 5.16a.

Fig. 5.16 shows the results and the evolution of the performance criteria for 10 repetitions of the optimization process. For comparison, we also display the IoU achieved with the baseline method introduced in Chapter 4, which does not require iterative optimization, but is restricted to rectangular shapes. We can see that the estimated number of edges converges seamlessly to 4, while the IoU goes to 1 for all the repetitions. Thus, our method achieves the performance of the baseline approach, despite not being provided with the ground-truth number of edges nor the rectangular constraint. Overall, the results demonstrate the potential of our approach for recovering polygonal shapes from acoustic echoes using a single mobile omnidirectional emitter/receiver device and beamforming.

Overall, we have provided an extension of the mapping approach first introduced in Chapter 4 to recover the geometry of unknown 2D polygonal structures from acoustic echos, and using a single omnidirectional mobile sensor. It combines delay-and-sum beamforming in the geometric



(a) Mapping setup and results for the experimental scenario.



(b) Evolution of performance criteria during the optimization.

Figure 5.16: Mapping results for the experimental setup. (a, left) shows the simulated sensor path (in black), the plate outline (discontinuous red) and estimated edges (green). The corresponding beamforming map is provided on the right, where the red rectangles indicate the retrieved edges. (b) shows the evolution of the loss, the number of estimated edges and the value of $\text{IoU}(\hat{\mathbf{M}}, \mathbf{M})$ for 10 repetitions of the optimizer. The black horizontal line on the IoU plot is the value achieved with the baseline method. The median values are in red.

boundary space and a simulated annealing optimizer to retrieve boundary estimates from the beamforming results. Contrary to previous works, our approach does not rely on solid hypotheses on echo detection, identification and association, and can recover, in principle, arbitrary polygonal shapes. In future works, the approach shall be evaluated in more realistic environments. Furthermore, the optimization efficiency shall be improved, so that it can be run online. It shall be integrated within a SLAM framework, as the sensor positions are to be estimated in practice, and its extension to recover 3D volumes may also be investigated.

5.6 Chapter summary

In this chapter, we essentially demonstrate the potential benefits of using combined representations along with UGWs measurements acquired by a mobile robot for mapping a metal structure. The core of this approach is the simultaneous recovery of a grid and a feature-based representation of a same metal plate environment. This allows maintaining both denseness and sparsity on the map belief, while leveraging the consistency properties between the maps.

First, we propose a Bayesian framework for building such joint maps. By leveraging the simplicity of DAS beamforming and the Bayesian binary filter, the proposed framework can fully and consistently map a rectangular metal plate. Compared to the feature-based only beamforming mapping method, the proposed approach is less sensitive to interference, it provides better uncertainty representation, while being robust to detection errors of the closest boundary.

In the second part of this chapter, we leverage such mapping framework for autonomous exploration, with the aim to automatically drive an autonomous robot for efficient mapping and coverage of a rectangular metal plate surface. The proposed algorithm, which relies on frontier exploration and on uncertainty assessment on the grid map belief, is evaluated both in simulation and in the real world. The results show that it is reasonable for the target application.

Finally, we propose to use a grid map of the inner surface of a plate to recover arbitrary polygonal shapes from ultrasonic echoes. Hence, we alleviate the need of the rectangular assumption on the feature-based map. The proposed approach, which introduces a regularization term on the cost criterion to retrieve line features on the DAS beamforming results, has proven to be efficient both in simulation and in a real-world setup to recover more arbitrary polygonal geometries, without prior knowledge of the number of line features to estimate. It is envisioned that the proposed frameworks can support the emergence of a true long-range robotic (and potentially automated) inspection operation.

Chapter 6

Graph-based Simultaneous Localization and Mapping

In this chapter, we propose a non-probabilistic GraphSLAM framework to solve the full SLAM problem using UGWs measurements. Contrary to online SLAM, which has been considered so far in this thesis, GraphSLAM can use all the measurements all at once to infer the trajectory and map estimates, usually leading to much more accurate results. We first elaborate an observation model that can leverage multi-path multi-order acoustic echoes for localization, while DAS beam-forming is used for mapping the plate boundaries. The two methods are subsequently merged within a nonlinear least squares optimizer to solve the full SLAM problem, without breaking the natural sparsity of the SLAM problem, in first approximation. We test our approach in real-world settings, on metal plates of varying size and material. The results demonstrate the improved localization accuracy of the multi-order echo observation model compared to a second model, from the previous chapters, that relies solely on first-order echoes. We also show that the entire robot trajectory and plate boundaries can be simultaneously and reliably recovered in two different scenarios, even when odometry noise induces severe drift.

6.1 Motivation and background

In general, a SLAM problem can come at two different tastes. When considering an online SLAM approach, the belief on the map and robot trajectory is often incrementally updated when a new measurement is made available. Solving the full SLAM problem, however, involves high-dimensional batch optimization to simultaneously recover the entire robot trajectory and map features based on all the observations. A conventional solution is provided by GraphSLAM [35]. The latter usually involves optimization over a sparse graph, where the unknown variables (robot positions and map features) are the nodes, and the edges between them represent information constraints that are induced either by odometry or measurements inputs. The objective is to find the unknown variables that minimize a cost criterion that is defined as a weighted sum of the individual constraint terms.

Unlike online filtering techniques, such as Kalman filtering or particle filtering [94], GraphSLAM can use the information contained in *all* the measurements to infer the target variables. This approach is typically known to lead to more accurate results, but at the cost of a higher computational effort. A promising prospect is that acoustic GraphSLAM, through a complete optimization over all the measurements, could provide by itself accurate robot localization, structure geometry mapping, and inspection results. Nevertheless, applying GraphSLAM techniques

to UGWs measurements is challenging due to the inherent complexity of ultrasonic measurements that present wave dispersion effects and multi-echoes characteristics.

To make one step towards long-range robotic inspection of large metal structures, we propose, in this chapter, a non-probabilistic GraphSLAM approach to simultaneously recover the robot trajectory and individual plate geometry from all odometry and acoustic measurements. Our objective is to enable the resolution of the full SLAM problem for accurate results, even in challenging acquisition conditions, as expected in a real-world inspection scenario. We introduce an observation model that can leverage the information contained in high-order reflections from the metal plate boundaries, and we rely on DAS beamforming for plate boundaries localization. The two methods are subsequently integrated within a Non-linear Least Squares (NLS) optimizer, which iteratively updates the trajectory and map estimates after linearization of a measurement likelihood cost criterion [48]. We show experimentally that the designed multi-order echo observation model leads to better localization accuracy when compared to the simpler first-order echo only model, introduced in the previous chapters. Furthermore, we demonstrate that the designed GraphSLAM approach can successfully solve the full SLAM problem in two different experimental setups, even when the drift induced by odometry noise is severe. Last, a study on the sensitivity of the SLAM results to the cost parameters is carried out.

6.1.1 Multi-order echoes signal processing

While traditional acoustic localization and mapping methods rely solely on the reception of the direct-path incident signal from the source, the additional integration of multi-path scattering signals is a good option for more accurate and more robust results. In UGW-based NDE, for example, the integration of multi-path signals can offer better defect detection capabilities [38, 105, 106]. Indeed, all multi-path echoes contain further information on the environment. Hence, using it (instead of treating it like noise) may lead to better results. However, due to the complexity of the signals, appropriate (and often more complex) models and algorithms are needed.

By doing a parallel with the former example, the localization methods presented earlier rely solely on detection of the first-order reflections on metal plate boundaries - or equivalently, they rely on direct incident signals originating from fictive first-order image sources. We now propose an observation model that can also leverage the high-order reflections in UGW measurements for more accurate and robust localization and mapping results on metal panels.

6.1.2 Related works on GraphSLAM

GraphSLAM is an approach that is now well studied [35]. In the literature, numerous GraphSLAM methods and solvers have been developed, with increasing computational efficiency [55, 21, 47, 46]. These techniques have gained such maturity that they can seamlessly adapt to various sensor modalities (Laser, computer vision, IMU...) and SLAM problems. The chore of state-of-the-art methods is to successively linearize the models to find a minimum of the cost criterion through non-linear least-squares. Furthermore, by leveraging the sparsity of the resulting Jacobian matrix, GraphSLAM can be used to recover increasingly large trajectories and maps, that can contain up to millions of features.

In the literature, acoustic SLAM using a single omnidirectional emitter/receiver device is the subject of various studies to enable, for example, simultaneous sensor localization and room geometry reconstruction from acoustic echoes due to reflections on the walls. Batch optimization methods are proposed in [52, 53] to solve the full SLAM problem. However, this is typically a

challenging task, as the acoustic data contain range-only information, and due to the multiple unlabeled echoes present in the measurements. For these reasons, these works are only tested in simulation, they don't explicitly address the problem of echo detection, and rely on essential assumptions, such as only first-order echoes are being observed, or the echo-map feature associations are flawlessly known. In [16], a graph-based SLAM approach that also integrates beamforming for DOA estimation is proposed for underwater localization with sound beacons. However, it requires a sensor array, and only direct incident signals are used.

6.2 A graph-based framework for ultrasonic SLAM

In the following, we present a novel GraphSLAM approach to solve the full SLAM problem by relying on ultrasonic measurements. We first introduce an observation model that can leverage information from high-order acoustic echoes for accurate localization. The latter approach is integrated, along with DAS beamforming, into a NLS optimizer.

6.2.1 Localization with a multi-order echo observation model

In a first step, we design an observation model that can leverage high-order reflections from UGWs measurements for localization only on a metal panel. Assuming a known plate geometry \mathbf{M} , we model, for a candidate sensor position \mathbf{x} , the expected envelope signal as a mixture of Gaussian functions:

$$\hat{e}(\mathbf{x}, r) = \sum_{\mathbf{p} \in \mathcal{I}(\mathbf{x}, \mathbf{M})} \mu_{\mathbf{p}} \cdot \exp \left\{ - \left(\frac{r - \|\mathbf{p} - \mathbf{x}\|}{\sigma} \right)^2 \right\}, \quad (6.1)$$

where $\mathcal{I}(\mathbf{x}, \mathbf{M})$ is the set of image source positions when the emitter is situated in \mathbf{x} , and the parameter value of σ can be chosen as the signal wavelength. As the amplitudes $\mu_{\mathbf{p}}$ may not be known a priori, we propose to estimate them by minimization of the mean-squared error between the modeled signal $\hat{e}(\mathbf{x}, r)$ and the envelope measurement $e_i(r)$: we first discretize the different variables: $\hat{\mathbf{e}}(\mathbf{x}) = [\hat{e}(\mathbf{x}, r_1), \hat{e}(\mathbf{x}, r_2), \dots, \hat{e}(\mathbf{x}, r_M)]^T$ and $\mathbf{e}_i = [e_i(r_1), e_i(r_2), \dots, e_i(r_M)]^T$, $M \in \mathbb{N}$, and define a "steering" vector function:

$$\mathbf{a}(\mathbf{x}, \mathbf{p}) = \left[\exp \left\{ - \left(\frac{r_1 - \|\mathbf{p} - \mathbf{x}\|}{\sigma} \right)^2 \right\}, \dots, \exp \left\{ - \left(\frac{r_M - \|\mathbf{p} - \mathbf{x}\|}{\sigma} \right)^2 \right\} \right]^T.$$

Let's call $N(\mathbf{x})$ the number of image sources \mathbf{p}_n , $n = 1..N(\mathbf{x})$, in $\mathcal{I}(\mathbf{x}, \mathbf{M})$ whose signal can be observed in \mathbf{e}_i , as they are such that $\|\mathbf{p}_n - \mathbf{x}\| < r_M$. After definition of the $M \times N(\mathbf{x})$ matrix:

$$\mathbf{A}(\mathbf{x}, \mathbf{M}) = [\mathbf{a}(\mathbf{x}, \mathbf{p}_1), \mathbf{a}(\mathbf{x}, \mathbf{p}_2), \dots, \mathbf{a}(\mathbf{x}, \mathbf{p}_{N(\mathbf{x})})],$$

the expected envelope measurement obtained at a candidate position \mathbf{x} is:

$$\hat{\mathbf{e}}(\mathbf{x}, \mathbf{M}, \mathbf{e}_i) = \mathbf{A}(\mathbf{x}, \mathbf{M}) \cdot \hat{\boldsymbol{\mu}}(\mathbf{x}, \mathbf{M}, \mathbf{e}_i),$$

where $\hat{\boldsymbol{\mu}}(\mathbf{x}, \mathbf{M}, \mathbf{e}_i)$ is the linear least square estimate:

$$\hat{\boldsymbol{\mu}}(\mathbf{x}, \mathbf{M}, \mathbf{e}_i) = [\mathbf{A}^T(\mathbf{x}, \mathbf{M})\mathbf{A}(\mathbf{x}, \mathbf{M})]^{-1} \mathbf{A}^T(\mathbf{x}, \mathbf{M}) \cdot \mathbf{e}_i. \quad (6.2)$$




Figure 6.1: Positions of the image sources (red dots) for a rectangular plate with dimensions $w \times h$ (blue), and a fixed emitter position $[x, y]$ (blue dot).

The positions of the image sources required to determine $\hat{\mathbf{e}}$ can be analytically derived for simple geometries based on the image source algorithm. For a rectangular panel with dimensions $w \times h$ and a source at $\mathbf{x} = [x, y]$ expressed in the plate frame, we have [2]:

$$\mathcal{I}(\mathbf{x}, \mathbf{M}) = \{[2nw \pm x, 2mh \pm y]^T / (n, m) \in \mathbb{Z}^2 / (0, 0)\}.$$

Fig. 6.1 illustrates the regular pattern formed by the image source positions, given a fixed real source. Fig. 6.2 illustrates how the envelope signal model can match the measurements. Fig. 6.2a shows an experimental ultrasonic measurement obtained on our aluminum plate, Fig. 6.2b is the corresponding envelope signal obtained after correlation with a wave propagation model expressed in the spatial domain. In Fig. 6.2c, the latter is superimposed with the expected envelope signal, which is calculated with Eq. (6.2) based on the ground-truth sensor position and plate size. It can be seen that there is a good match between the measured and expected envelopes, as there is an appropriate correspondence for all the positions of the local maxima.

Finally, to localize the robot on a metal plate given an envelope measurement \mathbf{e} , one may simply use the mean squared error between expected and measured envelope signals as a likelihood function of the sensor position:


$$l(\mathbf{x}|\mathbf{e}) \propto -\|\hat{\mathbf{e}}(\mathbf{x}, \mathbf{M}, \mathbf{e}) - \mathbf{e}\|^2.$$

The position of the maximum of this (highly non-linear) function yields a Maximum Likelihood (ML) robot position estimate. This will be further illustrated in the results section.


6.2.2 Mapping only with beamforming

Within the GraphSLAM framework, we propose to address the mapping problem using the classic DAS beamforming approach. As GraphSLAM necessitates the resolution of a highly non-linear problem, the integration of DAS beamforming may efficiently reduce the optimization effort. Indeed, beamforming may be used to integrate prior knowledge on how the map can be inferred given a candidate trajectory. Assuming a known $\mathbf{x}_{0:T}$, an estimate $\hat{\mathbf{M}}(\mathbf{x}_{0:T})$ of $\mathbf{M} = \{\mathbf{m}_l\}_{l=1..4} = \{r_l, \alpha_l\}_{l=1..4}$ is obtained by solving:


$$\hat{\mathbf{M}}(\mathbf{x}_{0:T}) = \arg \max_{\mathbf{M}} \sum_{(r, \alpha) \in \mathbf{M}} \mathcal{L}_{\mathbf{x}_{0:T}}(|x_i \cdot \cos \alpha + y_i \cdot \sin \alpha - r|),$$



(a) An experimental ultrasonic measurement.



(b) The resulting envelope of the correlation signal.



(c) The envelope signal shown in (b) superimposed with the expected envelope (in red) after matching the model in Eq. (6.1) with the envelope signal. A value $\sigma = 2\text{cm}$ is used.

Figure 6.2: Illustration of the envelope signal obtained from an experimental ultrasonic measurement acquired on a rectangular metal panel, and comparison with the expected envelope calculated with the matched Gaussian mixture model.




Figure 6.3: Graph representation of the (sparse) SLAM problem, where the variables to estimate are the positions \mathbf{x}_i and the map \mathbf{M} , the inputs are the odometry data \mathbf{u}_i and measurements \mathbf{z}_i . The constraints (induced by the external measurements) between the variables are represented by the graph edges, which are related to the l functions.

where $\mathcal{L}_{\mathbf{x}_{0:T}}$ is the usual DAS beamforming map:

$$\mathcal{L}_{\mathbf{x}_{0:T}}(r, \alpha) = \sum_{i=1}^T e_i(|x_i \cdot \cos \alpha + y_i \cdot \sin \alpha - r|),$$

and \mathbf{M} is restricted to be a rectangle. It should be noted that we restrict ourselves to the rectangular case for simplicity, even though we have proposed a methodology, in the previous chapter, to extend the geometry estimation to more arbitrary polygonal shapes.

6.2.3 The GraphSLAM cost criterion

In this section, we formulate the full SLAM problem using ultrasonic measurements as the minimization of a GraphSLAM criterion. GraphSLAM aims at finding a maximum a posteriori estimate of the full posterior distribution over the map and robot trajectory, given odometry and ultrasonic measurements. This amounts to minimizing the negative logarithm of the distribution that we model as a sum of non-linear mean-squared error terms:

$$J(\mathbf{x}_{0:T}, \mathbf{M}) = \sum_{i=0}^{T-1} \|\mathbf{x}_{i+1} - f_{\text{dyn}}(\mathbf{x}_i, \mathbf{u}_i)\|^2 + \lambda \cdot \sum_{i=0}^T \|\hat{\mathbf{e}}(\mathbf{x}_i, \mathbf{M}, \mathbf{e}_i) - \mathbf{e}_i\|^2, \quad (6.3)$$

where λ is a parameter that is used to balance the importance between odometry and ultrasonic measurements in the cost, f_{dyn} is the transition model, and \mathbf{u}_i are the odometry inputs. Overall, the introduced cost consists of two parts: the first sum accounts for constraints due to odometry data, and the second one accounts for the constraints due to measurements. Hence, the cost minimization simultaneously leads to a trajectory coherent with odometry inputs, and maximizes the likelihood of sensor presence based on accurate reconstruction of the envelope measurements. The corresponding graph of the SLAM problem is represented in Fig. 6.3. One may observe that it is a sparse graph. Due to the Markov property, the total number of edges grows linearly with the number of pose nodes, which leads to a sparsely connected graph. If the latter was fully connected, the increase of the number of edges would be quadratic with the number of nodes.




Figure 6.4: Graph representation of the SLAM problem after integration of DAS beamforming. The variables to estimate are the positions \mathbf{x}_i , and the map \mathbf{M} is hidden as it can directly be inferred from all \mathbf{x}_i using DAS beamforming. For simplicity of the graph representation, an edge is drawn to account for the existence of a constraint between two variables.

In our case-study, the minimization of J through search in the joint localization and map space can be quite inefficient. To reduce the problem complexity, the cost is only evaluated for candidate trajectories $\mathbf{x}_{0:T}$ and the corresponding map obtained by DAS beamforming $\hat{\mathbf{M}}(\mathbf{x}_{0:T})$. This enables to leverage the efficiency of DAS beamforming to integrate prior knowledge on how the map can be inferred given a trajectory. Thus, in the next, we will omit the map in the cost that will be reduced to:

$$J(\mathbf{x}_{0:T}) \triangleq J(\mathbf{x}_{0:T}, \hat{\mathbf{M}}(\mathbf{x}_{0:T})). \quad (6.4)$$

In this situation, however, the map \mathbf{M} becomes a "hidden" variable of the problem. As a result, the graph is no longer sparse, as the map introduces direct dependency between all the poses. As the graph of the SLAM problem has become fully-connected, as represented in Fig. 6.4, the resolution of the full SLAM problem can no longer be computationally efficient.

6.2.4 Sparse non-linear least squares minimization

Finding a minimum of Eq. (6.3) can be achieved through NLS minimization, which is conventional to solve GraphSLAM problems [35]. The principle is to iteratively update an initial guess of the {trajectory + map} variables based on a quadratic approximation of the non-linear cost.

Let's gather the unknowns in one variable $\mathbf{X} = (\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T)$. At each iteration $p + 1$ of the optimizer, the previous estimate $\mathbf{X}^{(p)}$ is updated with:

$$\mathbf{X}^{(p+1)} = \mathbf{X}^{(p)} + \Delta \mathbf{X}^{(p+1)}. \quad (6.5)$$

The increment $\Delta \mathbf{X}^{(p+1)}$ is found by minimization of the quadratic approximate of the cost criterion around $\mathbf{X}^{(p)}$. Based on the formulation of the cost in Eq. (6.4), the graph of the SLAM problem is no longer sparse due to the dependency between all the poses introduced by DAS beamforming. Thus, the determination of the increment may be computationally involving.

To reintroduce sparsity in the SLAM problem while leveraging DAS beamforming at the same time, we assume that the variation of the map induced by a "small" increment $\Delta \mathbf{X}^{(p+1)}$ in trajectory is *locally* negligible. Hence, only during calculation of the increment, the map can be assumed constant. This results in a sparse graph with the same structure as in Fig. 6.3, except that the map \mathbf{M} is not to be inferred anymore. As J is a sum of squared errors, under the assumption of constant map, the cost can be rewritten as:

$$J(\mathbf{X}) = \sum_{i,j} \mathbf{E}_{i,j}^T(\mathbf{x}_i, \mathbf{x}_j) \Sigma_{i,j} \mathbf{E}_{i,j}(\mathbf{x}_i, \mathbf{x}_j) \triangleq \mathbf{E}^T(\mathbf{X}) \Sigma \mathbf{E}(\mathbf{X}), \quad (6.6)$$

where the $\mathbf{E}_{i,j}$ are vector functions that assess how much the constraint relating \mathbf{x}_i and \mathbf{x}_j is satisfied. \mathbf{E} is a stacked vector containing all the $\mathbf{E}_{i,j}$, and $\mathbf{\Sigma}$ is a block-diagonal matrix with block elements $\mathbf{\Sigma}_{i,j}$. The $\mathbf{\Sigma}_{i,j}$ are matrix parameters of the cost function, which can be deduced from Eq. (6.4). Due to sparsity of the SLAM problem, there is only a few indices (i, j) for which it is non-zero. In the case of an odometry constraint, $\mathbf{\Sigma}_{i,j} \neq 0$ when $j = i + 1$, leading to:

$$\mathbf{E}_{i,i+1}(\mathbf{x}_i, \mathbf{x}_{i+1}) = \mathbf{x}_{i+1} - f_{\text{dyn}}(\mathbf{x}_i, \mathbf{u}_i).$$

For a constraint related to an ultrasonic measurement, we have $\mathbf{\Sigma}_{i,j} \neq 0$ when $i = j$, and:

$$\mathbf{E}_{i,i}(\mathbf{x}_i, \mathbf{x}_i) = \hat{\mathbf{e}}(\mathbf{x}, \hat{\mathbf{M}}, \mathbf{e}_i) - \mathbf{e}_i.$$

Note that there is no dependency in the map, as $\hat{\mathbf{M}}$ is assumed constant. To determine $\Delta \mathbf{X}^{(p+1)}$ at iteration $p+1$ of the NLS optimizer, the full vector function \mathbf{E} is linearized around the current trajectory candidate, and using a constant map $\hat{\mathbf{M}}^{(p)}$:

$$\mathbf{E}(\mathbf{X}^{(p)} + \Delta \mathbf{X}) \approx \mathbf{E}^{(p)} + \mathbf{J}^{(p)} \Delta \mathbf{X}, \quad (6.7)$$

where $\mathbf{J}^{(p)} = \mathbf{J}(\mathbf{X}^{(p)})$ is the Jacobian matrix of \mathbf{E} in $\mathbf{X} = \mathbf{X}^{(p)}$, and $\mathbf{E}^{(p)} = \mathbf{E}(\mathbf{X}^{(p)})$ for notation simplicity. One should note that, due to the complexity of the cost, the Jacobian matrices $\mathbf{J}^{(p)}$ should be numerically approximated. Substituting Eq. (6.7) in Eq. (6.6) yields:

$$\begin{aligned} J(\mathbf{X}^{(p)} + \Delta \mathbf{X}) &\approx \left[\mathbf{E}^{(p)} + \mathbf{J}^{(p)} \Delta \mathbf{X} \right]^T \mathbf{\Sigma} \left[\mathbf{E}^{(p)} + \mathbf{J}^{(p)} \Delta \mathbf{X} \right], \\ &\approx \left(\mathbf{E}^{(p)} \right)^T \mathbf{\Sigma} \mathbf{E}^{(p)} + 2 \left(\mathbf{E}^{(p)} \right)^T \mathbf{\Sigma} \mathbf{J}^{(p)} \Delta \mathbf{X} + \Delta \mathbf{X}^T \left(\mathbf{J}^{(p)} \right)^T \mathbf{\Sigma} \mathbf{J}^{(p)} \Delta \mathbf{X}, \\ &\approx C + 2\mathbf{B}^T \Delta \mathbf{X} + \Delta \mathbf{X}^T \mathbf{H} \Delta \mathbf{X}, \end{aligned}$$

which is the quadratic approximation of the cost. The first-order optimality condition is:

$$\frac{\partial}{\partial \Delta \mathbf{X}} J(\mathbf{X}^{(p)} + \Delta \mathbf{X}) \approx 2\mathbf{B} + 2\mathbf{H} \Delta \mathbf{X} = 0.$$

Hence, the increment $\Delta \mathbf{X}^{(p+1)}$ is found by solving the linear system:

$$\mathbf{H} \Delta \mathbf{X} = -\mathbf{B}. \quad (6.8)$$

Due to multiple zero components of $\mathbf{\Omega}$, \mathbf{H} is a sparse matrix. This property can be leveraged by conventional linear numerical solvers through Cholesky factorization of \mathbf{H} , for efficient resolution of Eq. (6.8) [35]. After application of Eq. (6.5) to update the trajectory, the map is updated with DAS beamforming: $\hat{\mathbf{M}}^{(p+1)} = \hat{\mathbf{M}}(\mathbf{X}^{(p+1)})$. Then, the iteration process is repeated until a termination criterion is fulfilled.

6.3 Experimental results

We assess the efficiency of our GraphSLAM approach on two sets of acoustic measurements acquired on our two rectangular metal panels. We use the dataset acquired on the 600x450x6mm aluminum plate for Scenario 1, and the UGWs measurements acquired on the 1700x1000x6mm steel plate for Scenario 2.

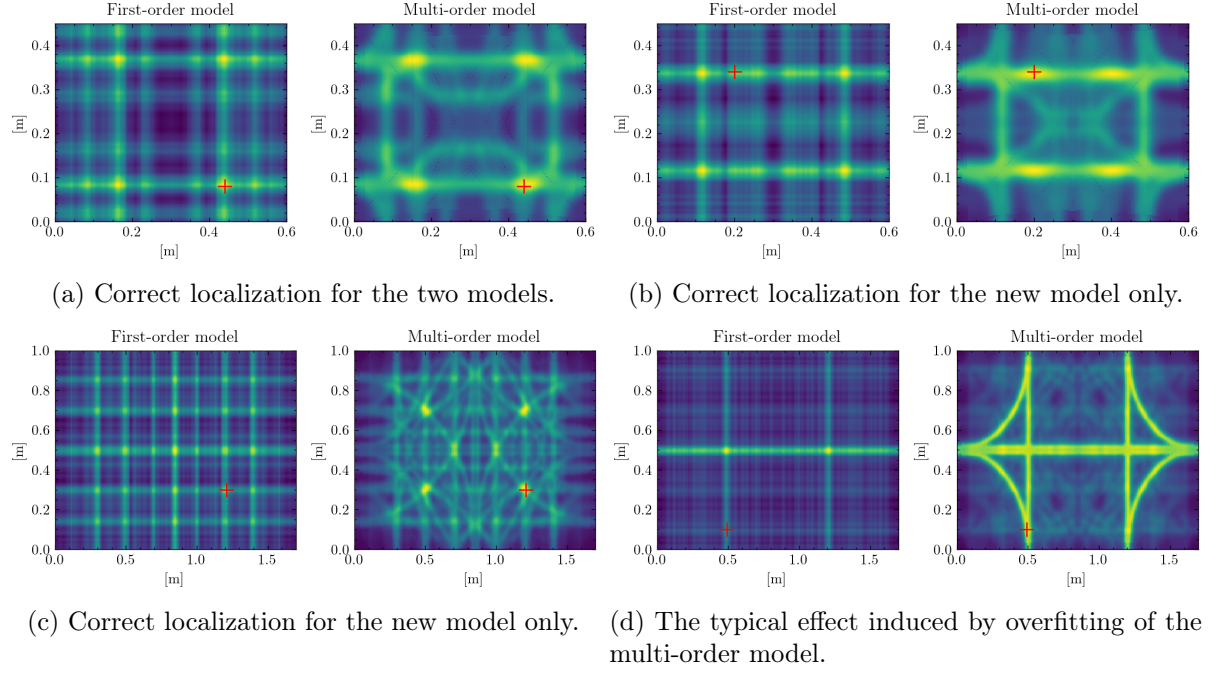


Figure 6.5: Likelihood maps of sensor presence obtained using measurements acquired on the aluminum plate (a), (b), and on the steel plate (c), (d). The true sensor position is indicated by the red mark. The outline of each plot accounts for the plate boundaries, and the color maps only spans the range from the minimum to the maximum likelihood value.

6.3.1 Comparison of the observation models performance

To assess the benefit of using a multi-order echo observation model for localization, we first compare the localization accuracy achieved both with the proposed multi-order echo model, and the simpler model used in the previous chapters, that only relies on first-order reflections. We assess the likelihood of sensor presence at position \mathbf{x} for each model with:

$$l_{1^{\text{st}}\text{order}}(\mathbf{x}|\mathbf{e}_i) = \sum_{(r,\alpha) \in \mathbf{M}} e_i(|x \cdot \cos \alpha + y \cdot \sin \alpha - r|)$$

for the first-order model, and:

$$l_{\text{multi-order}}(\mathbf{x}|\mathbf{e}_i) = -||\hat{\mathbf{e}}(\mathbf{x}, \mathbf{M}, \mathbf{e}_i) - \mathbf{e}_i||^2$$

for the multi-order echo model. Likelihood maps are determined for each model, using different envelope measurements e_i , and based on a known plate geometry \mathbf{M} . A likelihood maximum is expected at the true sensor position. We show, in Fig. 6.5, several instances of likelihood maps obtained for each model, and that illustrate the pattern tendencies based on measurements acquired on the two metal panels. It can be noted that the likelihood maps present the same two axis symmetries than the plates due to the ambiguity of the echo measurements, which cannot be disambiguated without an external localization system.

We first qualitatively assess the performance of the observation models. Fig. 6.5a depicts a scenario where the two models point to the correct sensor position. In Fig. 6.5b, only the first-order model does not point to the correct sensor position. This is due to interference, that can wrongly attribute one energetic reflection in the envelope signal to reflections from a different




Figure 6.6: Average localization errors calculated using the two considered models. The position error is averaged over all the measurements acquired on the aluminium plate (top) and steel plate (bottom) with respect to γ , yielding the solid lines. The upper and lower bounds of the colored areas are situated respectively above and below one standard deviation from the average.

boundary, leading to incorrect localization. Furthermore, highly energetic high-order reflections also contribute to localization errors because the model treats them as first-order ones. With the multi-order echo model, however, the likelihood maximum can be identified at the ground truth sensor position. The same observations hold for Fig. 6.5c, which illustrates the likelihood maps for a measurement acquired on the large steel plate. This illustrates that the multi-order approach is better suited for localization. Fig. 6.5d illustrates the typical detrimental effect of overfitting with the multi-echo model, which occurs when one echo level is relatively high compared to the others in the envelope measurement. Due to the estimation of the coefficients $\mu_{\mathbf{p}}$ in Eq. (6.1) by matching expected and measured envelope signals, the former will present substantial similarity with the measurement for any position for which the echo is expected (either as a first or a high-order echo). This results in higher localization ambiguity. Yet, this phenomenon is observable in the dataset for only a very limited number of measurements. We will see later if it detrimentally affects the GraphSLAM results.

To compare the model’s performances quantitatively, we assess the average localization error calculated over the whole measurement sets, for each of the two metal plates. The initial likelihood maps are first re-scaled to values between 0 and 1, and are passed through an exponential mapping function using a multiplicative scaling parameter γ . Next, the average localization errors are assessed, after considering the resulting maps as (non-normalized) distributions over the sensor position on the plate, with:

$$\Delta_i \mathbf{x}_{1^{\text{st}}\text{-order}}(\gamma) = \frac{\sum_{\mathbf{x}} \|\mathbf{x} - \mathbf{x}_i^{\text{true}}\| \cdot \exp \gamma l_{1^{\text{st}}\text{-order}}^{\text{scaled}}(\mathbf{x}|\mathbf{e}_i)}{\sum_{\mathbf{x}} \exp \gamma l_{1^{\text{st}}\text{-order}}^{\text{scaled}}(\mathbf{x}|\mathbf{e}_i)},$$

$$\Delta_i \mathbf{x}_{\text{multi-order}}(\gamma) = \frac{\sum_{\mathbf{x}} \|\mathbf{x} - \mathbf{x}_i^{\text{true}}\| \cdot \exp \gamma l_{\text{multi-order}}^{\text{scaled}}(\mathbf{x}|\mathbf{e}_i)}{\sum_{\mathbf{x}} \exp \gamma l_{\text{multi-order}}^{\text{scaled}}(\mathbf{x}|\mathbf{e}_i)},$$

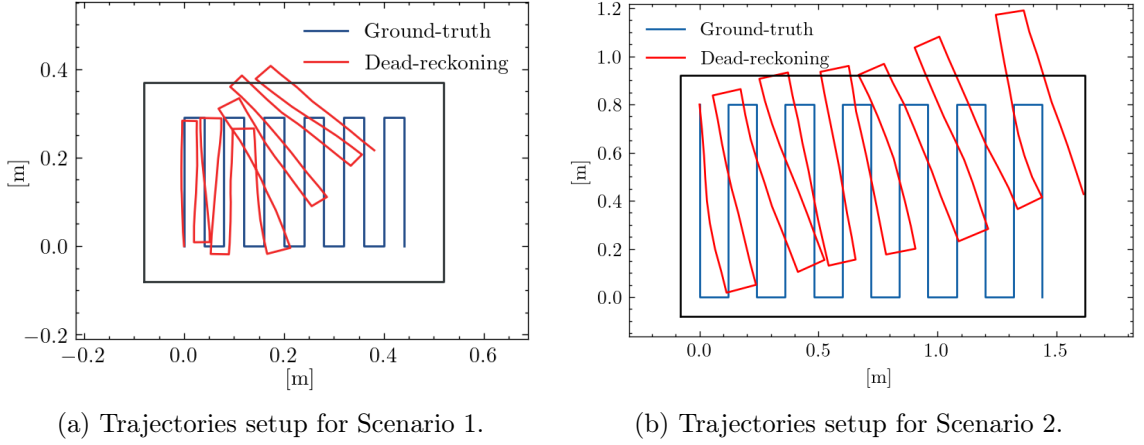


Figure 6.7: Ground truth trajectories and dead-reckoning trajectories used for Scenario 1 (a) and Scenario 2 (b). The black rectangle represents the actual outline of the metal plate.

where $\mathbf{x}_i^{\text{true}}$ is the ground truth sensor position, and γ is a positive scaling parameter that balances the ambiguity level due to multiple likelihood maxima. Indeed, for a low γ , different likelihood maxima can contribute to the localization error, whereas the error reduces to that achieved by the maximum-a-posteriori position estimate when γ goes to infinity.

Fig. 6.6 shows the evolution of the average localization errors and standard deviations calculated for the two models, and w.r.t. the scaling parameter γ . The errors are evaluated for both the data acquired on the aluminium plate and the data acquired on the steel plate. It can be observed that, in the two scenarios, and no matter the value of γ , the multi-order model outperforms the baseline model in terms of localization accuracy.

Overall, these results support the benefit of using our approach that can leverage multi-order echoes for localization, compared to first-order echo only models. The following subsection will illustrate that, when integrated within the proposed GraphSLAM approach, the multi-order model can naturally lead to more accurate results than those presented in Fig. 6.6.

6.3.2 Experimental localization and mapping results

We now assess the performance of the designed GraphSLAM approach on the experimental data for solving the full SLAM problem. We consider lawn-mower transects as simulated robot trajectories, and use the theoretic displacement between two measurements as noise-free odometry. The resulting linear and rotational displacements are corrupted with additive Gaussian noise: $\mathbf{u}_i = [\tilde{\Delta}_{r,i}, \tilde{\Delta}_{\theta,i}]^T$, with $\tilde{\Delta}_{r,i} \sim \mathcal{N}(\Delta_{r,i}, 0.006^2)$ and $\tilde{\Delta}_{\theta,i} \sim \mathcal{N}(\Delta_{\theta,i}, 0.04^2)$. Fig. 6.7 shows the ground truth trajectories used for the two considered scenarios, along with the dead-reckoning trajectory obtained based on the generated noisy odometry data. One may observe, especially for Scenario 1, that the drift induced by odometry noise is relatively more important than the drift in the experiments of Chapter 4.

To consider a baseline for comparison, we run the FastSLAM algorithm presented in Chapter 4, in the situation corresponding to Scenario 1. The localization and mapping results obtained at different steps of the simulated robot trajectory are presented in Fig. 6.8. One can observe the overall poor localization and mapping results. Due to its inability to leverage future measurements to adjust incorrect past position estimates (and particle sampling following an incorrect map), FastSLAM is unable to recover the trajectory and metal plate geometry in that scenario

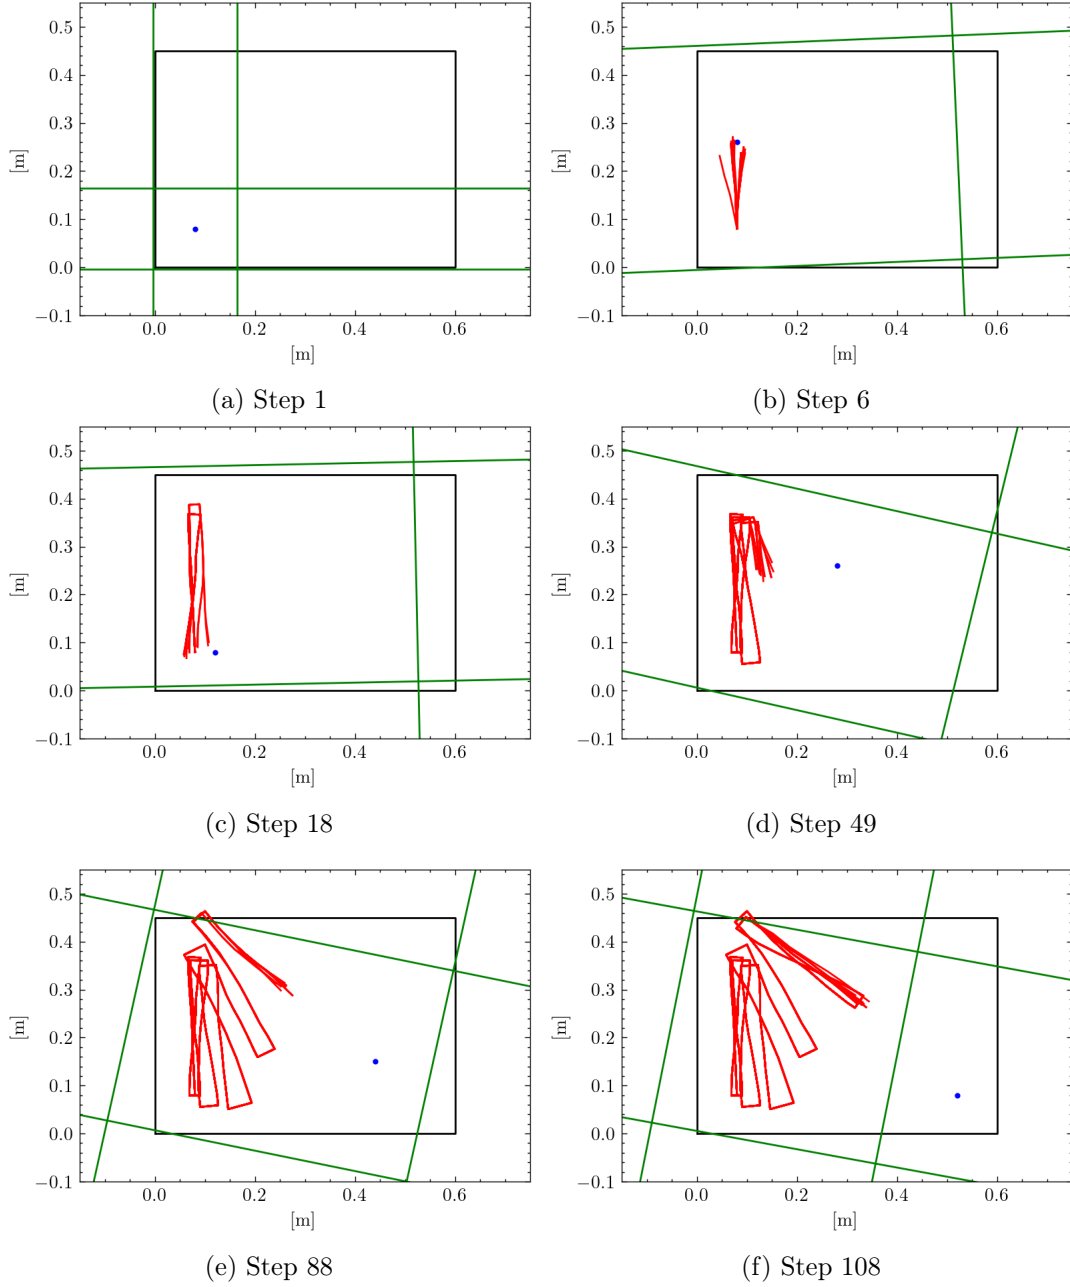


Figure 6.8: Localization and mapping results obtained with FastSLAM for the Scenario 1. The figure represents the trajectories estimated by all the particles (red lines), along with the map retrieved by the particle with the highest weight (green lines) during Steps 1, 6, 18, 49, 88 and 108 of the simulated lawn-mowing trajectory. The true outline of the plate and true sensor positions correspond to the black rectangle and blue dot, respectively.

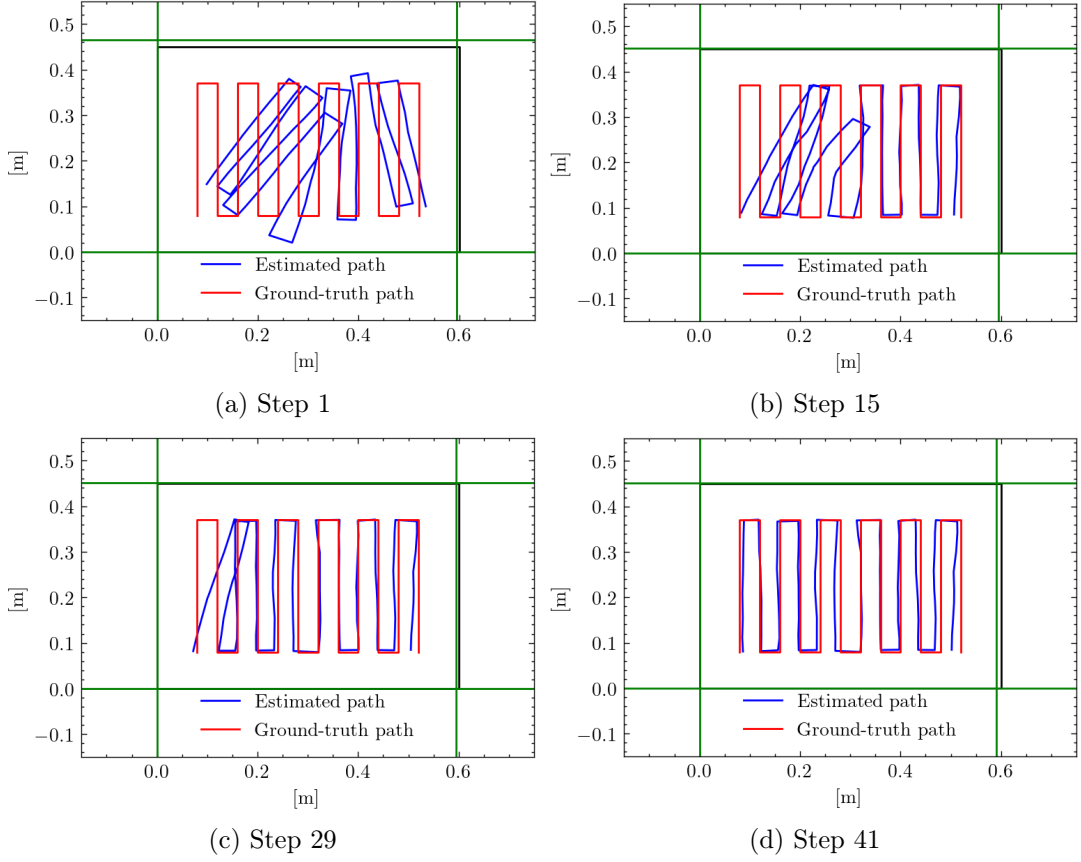


Figure 6.9: Representation of the SLAM results during iteration Steps 1, 15, 29, and 41 of the NLS optimization process for Scenario 1. The true outline of the plate is represented by the black rectangle. For better comparison of the reconstructed trajectory with the ground truth, they are expressed in the estimated (resp. ground truth) plate frame, while the origin of the estimated plate geometry is fixed at $[0, 0]$.

where the odometry drift is severe.

We now run our GraphSLAM approach on the data of Scenario 1. We use the implementation of the NLS optimizer³ provided by the Scipy library, on Python. To leverage the sparsity of the GraphSLAM problem, we also provide the structure of the Jacobian matrix (*i.e* the positions of the non-zero components). This greatly improves the computation of the Jacobian matrix following the procedure in [20]. The matrix is numerically approximated using the default finite difference "2-points" method. The optimizer is initialized with the dead-reckoning trajectory, and the map is initialized with the plate geometry obtained with DAS beamforming. We then run the optimizer, considering a value of $\lambda = 0.03^2$ for the cost function parameter, and $\sigma = 2$ cm for the envelope signal model parameter. The iterative optimization terminates when either the relative change of the cost function, the relative change of the independent variables, or the change in the gradient vector become inferior to their corresponding (default) value.

With our implementation, the average execution time for one NLS iteration is 2.30 seconds. As several tens of iterations might be necessary to reach a satisfying solution, the approach

³Documentation available at: https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.least_squares.html.




Figure 6.10: Evolution of the loss (left), along with the localization and mapping errors (right) during the NLS optimization process for Scenario 1. The loss evaluated at ground truth is represented as a horizontal line.




Figure 6.11: Evolution of the loss (left), along with the localization and mapping errors (right) during the NLS optimization process for Scenario 2. The loss evaluated at ground truth is represented as a horizontal line.

cannot be run online. Yet, one may expect a total execution time of 1 or 2 minutes, which appears, for now, to be reasonable for an offline SLAM approach. Fig. 6.9 illustrates the SLAM results at different steps of the optimization process. It can be observed in Fig. 6.9a that initially, the plate size estimated via DAS beamforming is accurate, but the localization error is seemingly important. Step by step, the estimated trajectory converges to the ground truth one. During Step 41, both the map and trajectory have been fully recovered.

In Fig. 6.10, we show the evolution, during 50 iterations of the NLS optimization process, of the mapping error, as well as the evolution of the localization error. The former is defined as the average error on plate dimensions, while the latter is evaluated as the mean distance between ground truth and estimated trajectories expressed in the ground truth and estimated plate frames, respectively. One may first observe that, after convergence of the optimizer, a loss value lower than the loss evaluated at ground truth is achieved, which illustrates the efficiency of our implementation of NLS. One may also note that the mapping error does not significantly vary, it eventually reaches an error of 1 cm. This is due to the efficiency of DAS beamforming for recovering the plate size despite important localization error. Besides, the localization error, which is 6 cm initially, is reduced to 1 cm after convergence.

To determine if our approach still performs well on measurements acquired on a larger plate surface, and for a different plate material, we perform the same experiment using the data acquired on the steel plate (Scenario 2). Fig. 6.11 shows the evolution of the loss along with the localization and mapping errors during the NLS optimization. It can still be observed that, after convergence, the loss is lower than that achieved at ground truth, which illustrates the




Figure 6.12: Representation of the SLAM results during iteration step of the NLS optimization process for Scenario 2. The true outline of the plate is represented by the black rectangle. For better comparison of the reconstructed trajectory with the ground truth, they are expressed in the estimated (resp. ground truth) plate frame, while the origin of the estimated plate geometry is fixed at $[0, 0]$.

efficiency of the optimization process. Furthermore, the average localization and mapping errors drop from 3 and 5 cm to around 2 cm. One may note that the decrease of the mapping error is more noticeable in this second scenario. Fig. 6.12 shows the SLAM results in this scenario, where one can observe the satisfying matching between estimated map and trajectory with the ground truths. Overall, these results demonstrate the potential of our GraphSLAM approach that leverages a multi-order echo model and a NLS optimizer to achieve SLAM with ultrasonic guided waves on a metal plate.

6.3.3 A study on the sensitivity of GraphSLAM to the loss parameter

In this section, we evaluate the sensitivity of the SLAM results to the cost function parameter. To this end, we run the NLS optimizer on the data from Scenario 1 for different values $\lambda \in [0.001^2, 0.005^2, 0.02^2, 0.05^2, 0.1^2]$. In Fig 6.13, we display the evolution of the difference between the loss achieved by the optimizer, and the loss at ground truth, along with the evolution of the localization and mapping errors, for each value of λ .

One may first notice that for $\lambda = 0.05^2$ and $\lambda = 0.1^2$, the cost difference is not negative after convergence, which results in an inefficient optimization. The optimizer, which iteratively approximates the nonlinear cost criterion, becomes stuck in a local minimum, resulting in poor localization results. A value of $\lambda = 0.02^2$ yields the best localization performance. For both $\lambda = 0.001^2$ and $\lambda = 0.005^2$, the cost difference is already negative at the start of the optimization process, which indicates that the confidence in the ultrasonic measurements is underestimated. In the first case, however, there is no sensible decrease of the localization error, whereas in the second one, the localization error is less than 1 cm, after convergence. There is little variation in the mapping error, which remains lower than 1 cm in all the scenarios. This can be explained by the robustness of DAS beamforming to localization error, especially when the trajectory is initialized with dead-reckoning path.

These observations illustrate the sensitivity of the NLS approach to the cost parameter value, which is critical to accurate SLAM results. As it may not be clear which value to choose




Figure 6.13: Evolution of loss along with the localization and mapping errors during the NLS optimization in Scenario 1, and for different cost parameters λ . The top left graph shows the evolution of the difference between the achieved loss and the loss at ground truth. A close view of the same graph is provided on the bottom left plot. The middle plot shows the evolution of the localization errors, while the mapping errors are provided on the right plot.

beforehand, one may experiment with different λ , and determine a posteriori an interval that leads to a consistent and stable trajectory (likely to be close to $\lambda \in [0.01^2, 0.05^2]$ in our scenario).

6.3.4 Discussions

The approach we proposed to solve the full SLAM problem using ultrasonic measurements can accommodate multiple high-order reflections without performing explicit echo extraction and association, which are known to be challenging tasks.

Nevertheless, the proposed approach has several drawbacks: the proposed GraphSLAM approach can lead to an increase in performance when compared to FastSLAM. Yet, the experiments have illustrated its sensitivity to the parameter value of the cost criterion, which needs to be carefully tuned for accurate results. Besides, the uncertainty estimation on the SLAM result has not been investigated in the present study, whereas indicators on the confidence one can have on the solution reached by the optimizer can be crucial, especially in noisy environments. Also, the proposed approach is computationally involved. In the experiments, the sensor positions and map features are in relatively low numbers, but the method cannot be run online. In contrast, state-of-the-art GraphSLAM methods, which are computationally efficient, can cope with tens of thousands of poses and up to millions of map features. Finally, our approach, as presented here, is limited to plates with rectangular geometries. This is not an issue for an inspection operation on ship hulls or storage tanks, as they are almost entirely made of rectangular panels. However, the approach could be extended to work with more general geometries by relying on the more general mapping framework presented in the previous chapter.

To further assess the viability of our approach for the target application of long-range robotic inspection, the former shall be evaluated in more realistic conditions. Indeed, a lower signal-to-noise ratio is expected on an actual structure due to larger plates, background environmental noise, the low reflectivity of weld joints, or the presence of coating. Apart from application to robotic inspection, we envision that our approach can be used for other acoustic SLAM problems, especially in environments that present 3-dimensional reflective surfaces, through adaptation of the distance function and the map feature space.

6.4 Chapter summary

In this chapter, we elaborate a non-probabilistic graph-based framework to solve the full SLAM problem for an inspection robot relying on ultrasonic guided waves on a metal structure. We introduce an observation model that can leverage high-order echoes within the measurements for localization, and we rely on DAS beamforming to localize the plate boundaries. We integrate these elements within a custom implementation of a NLS framework to optimize a GraphSLAM cost criterion. Experimental results performed in a laboratory environment show that the designed observation model yields better localization accuracy than our initial first-order echo model, and it can yield more accurate results, close to sub-centimeter precision, when compared to FastSLAM, even in the presence of severe odometry drift.

In future works, the computational efficiency of the approach shall be improved to work with more significant trajectories and map features, and the uncertainty assessment on the optimizer result must be investigated. Also, it shall be tested in more realistic conditions to assess its viability for inspection on actual structures, in outdoor environments. Finally, other map features such as defects should be included in the map for a complete inspection task.

Chapter 7

Learning the propagation model for Lamb wave-based mapping

The inspection of sizeable plate-based metal structures such as storage tanks or marine vessel hulls is a significant stake in the industry, which necessitates reliable and time-efficient solutions. Although Lamb waves have been identified as a promising solution for long-range non-destructive testing, and despite the substantial progress made in autonomous navigation and environment sensing, a Lamb-wave-based robotic system for extensive structure monitoring is still lacking. Following our previous work on ultrasonic Simultaneous Localization and Mapping, we introduce a method to achieve plate geometry inference without prior knowledge of the material propagation properties, which may be lacking during a practical inspection task in challenging and outdoor environments. Our approach combines focalization to adjust the propagation model parameters and DAS beamforming to infer the plate boundaries location by relying directly on acoustic measurements acquired along the mobile unit trajectory. For each candidate model, the focusing ability of the corresponding beamformer is assessed over high-pass filtered beamforming maps to further improve the robustness of the plate geometry estimates. We then recover the optimal space-domain beamformer through a simulated annealing optimization process. We evaluate our method on three sets of experimental data acquired in different conditions, and show that accurate plate geometry inference can be achieved without any prior propagation model. Finally, the results show that the optimal beamformer outperforms the beamformer resulting from the predetermined propagation model in non-nominal acquisition conditions.

7.1 Introduction and background

The development of an accurate ultrasound-based localization and mapping system is key to the emergence of a true long-range robotic inspection operation. Indeed, on a large metal structure, a moving robot, equipped with an omnidirectional emitter/receiver ultrasonic sensor, may be able to detect structure-bound landmarks, such as weld joints or stiffeners, and subsequently use them for accurate localization, within a SLAM framework. After determination of the structural elements along with the measurement positions, acoustic tomography could be achieved to identify small acoustic scatters that could be potential defects.

To make one step towards such a challenging application, we have first proposed a SLAM framework for localization and metal plate geometry reconstruction. The core of our methodology is DAS beamforming integrated in a FastSLAM algorithm (for the online estimation problem), or within a GraphSLAM framework (for the offline batch estimation problem). The experimental

results demonstrated the efficiency of the proposed methods on two different metal panels, in a laboratory environment. Yet, a key element is necessary for the successful achievement of ultrasonic SLAM: the prior knowledge of an accurate wave propagation model.

Prior knowledge of the physical properties of isotropic metal structures may be sufficient for acoustic localization and mapping in well controlled environments. However, the hypothesis that the propagation model is known a priori may not be realistic for a practical inspection task on a large metal structure due to a wide variety of external perturbations that can significantly affect the acoustic signals, and due to a lack of knowledge of the structure physical condition (which is the reason why it is being inspected). Potential perturbations may include, but may not be limited to, temperature variations, which can substantially affect the acoustic measurements [12], the pressure due to the neighboring metal plates which are welded altogether, moisture, the varying thickness of the coating, or the effect of the coupling (namely water in the case of acoustic inspection with a robotic system). Consequently, adaptive methods that automatically calibrate the propagation model by relying directly on data acquired on the field may be necessary to achieve accurate localization and mapping results, without any human intervention.

In the next, we propose an approach that combines acoustic focalization and DAS beamforming to map a plate-like structure by a mobile unit, so that accurate mapping results can be achieved during an inspection task without prior knowledge of the propagation properties of the material. The proposed approach combines focalization to adjust the parameters of propagation models that are derived from the Helmholtz equation, and DAS beamforming to localize the plate edges. The focalization capability of a candidate beamformer is assessed over high-pass filtered beamforming maps, so that the detrimental effects of interference and high-order reflections are lowered. A Simulated Annealing (SA) optimization process is then implemented so that the optimal beamformer can be recovered based directly on data. Results acquired on experimental data in three different scenarios show that the proposed approach is effective. Furthermore, mapping the plate geometry with the optimal beamformer is found to be more efficient than mapping with a predetermined propagation model in non-nominal acquisition conditions. For the sake of simplicity, the localization problem is not addressed here, but the previous chapters provide indications on how localization could be considered as well.

7.1.1 Related works

Recent works in the literature either address the problem of extracting the dispersion characteristics of materials from ultrasonic data [104, 13, 18, 14], or investigate acoustic localization techniques that can also recover the propagation properties [83, 56, 103]. However, simultaneous acoustic localization and propagation characterization from pulse-echo ultrasonic measurements acquired by a mobile unit has not been thoroughly studied. In the literature of beamforming, adaptive methods have been elaborated to compensate for inaccurate knowledge of the properties of the propagation media for underwater acoustic source localization purposes [81], or more generally, for Matched Field Processing [7]. Beamforming has been combined with focalization [19], an approach that considers the propagation environment as an acoustic lens, and which seeks, through an optimization process, to adjust the propagation parameters so that the focalization capability of beamforming is maximized. It has also been studied for joint acoustic signal separation and source localization [57]. The potential of focalization has been successfully demonstrated in simulation for underwater acoustic localization, and may provide an interesting approach for Lamb waves. Yet, this has not yet been demonstrated.

7.2 Proposed method

In what follows, we propose an approach for mapping the geometry of a metal panel without prior knowledge of the propagation model. We first briefly recall the propagation models based on approximate solutions of the Helmholtz equation. Next, we summarize the mapping strategy to localize the plate boundaries based on a known propagation model and DAS beamforming. These elements will be the basis of an optimal beamforming process, that can adjust the propagation model parameters through maximization of the focusing ability of beamforming.

7.2.1 Lamb wave-based mapping

Following our general setup, we are considering a mobile platform equipped with an emitter/receiver pair of piezoelectric transducers nearly co-located, and moving on a metal surface. All along a trajectory $\mathbf{x}_{0:T}$ (assumed to be flawlessly known, as we are only interested in the mapping problem), an excitation signal s is used to generate a Lamb wave in the material. Ultrasonic measurements $z_i(t)$, which contain the reflections on the plate boundaries, are subsequently collected. As for all the chapters, it will be assumed that the excitation signal is chosen adequately so that the A0 mode is predominant (while the propagation of S0 is negligible). The objective is to infer, by relying on an isotropic propagation model of the Lamb wave, the plate geometry and the sensors' acquisition positions in the plate frame.

Following the approach used so far, a wave propagation model is first determined based on prior knowledge of the plate material. The latter is recalled here for convenience. The wave dispersion relation $k(\omega)$ for the A0 mode is calculated by numerically solving the corresponding Rayleigh Lamb equation:

$$\frac{\tanh qh}{\tanh ph} = -\frac{(k^2 - q^2)^2}{4k^2qp},$$

where h denotes the half-thickness of the plate, $k = 2\pi/\lambda$ is the wavenumber, λ is the wavelength, and p and q are defined with:

$$p^2 = \frac{\omega^2}{c_L^2} - k^2; \quad q^2 = \frac{\omega^2}{c_T^2} - k^2.$$

In the above equation, c_L and c_T are respectively the longitudinal and transverse velocities related to the plate material. Their numerical values can be found in lookup tables [96]. Next, based on the wave dispersion relation, a model of the transfer function is derived with:

$$\hat{g}(r, \omega) \approx e^{-jk(\omega)r} / \sqrt{k(\omega)r}, \quad (7.1)$$

where r is the propagation distance. To determine a plate geometry estimate from all the measurements, we first determine the envelope signals e_i :

$$e_i(r) = |z'_i(r) + j\mathcal{H}(z'_i)(r)|,$$

where z'_i is obtained through correlation of the measurement with the wave propagation model:

$$z'_i(r) = \frac{\langle z_i(t), \hat{z}(r, t) \rangle}{\sqrt{\langle z_i(t), z_i(t) \rangle \langle \hat{z}(r, t), \hat{z}(r, t) \rangle}},$$

where $\hat{z}(r, t) = \hat{g}(2r, t) * s(t)$ is the expected signal for the incident wave reflecting at a distance r from the transducers. Next, an estimate of the map $\mathbf{M} = \{r_l, \alpha_l\}_{l=1..4}$ is obtained by applying

DAS beamforming to the envelope measurements:

$$\mathcal{L}_T(r, \alpha) = \sum_{i=0}^T e_i(|x_i \cdot \cos \alpha + y_i \cdot \sin \alpha - r|), \quad (7.2)$$

and subsequently solving the following optimization problem:

$$\hat{\mathbf{M}} = \arg \max_{\mathbf{M}} \mathcal{L}_T(\mathbf{M}) = \arg \max_{\mathbf{M}} \sum_{l=1}^4 \mathcal{L}_T(r_l, \alpha_l) \quad (7.3)$$

where \mathbf{M} is restricted to be a rectangle. We have proposed to solve the problem efficiently by taking that constraint into account. First, we determine the most likely line boundary with:

$$(\hat{r}_1, \hat{\alpha}_1) = \arg \max_{r, \alpha} \mathcal{L}_T(r, \alpha).$$

Next, assuming that $\hat{\alpha}_1$ provides the most reliable estimation of the plate orientation w.r.t. the robot, the determination of the other lines reduces to solving independent and straightforward one-dimensional optimization problems:

$$\hat{\alpha}_l = \hat{\alpha}_1 + \frac{\pi(l-1)}{2} ; \quad \hat{r}_l = \arg \max_r \mathcal{L}_T(r, \hat{\alpha}_l)$$

for $l = 2, 3, 4$. The experiments from the previous chapters have shown that this approach can give accurate results in laboratory conditions. Yet, one of the major limitations of this approach is that it relies on prior knowledge of an accurate propagation model \hat{g} to obtain precise results. However, such a hypothesis may not be realistic for a practical inspection task in challenging outdoor environments, where the structure state is truly unknown.

7.2.2 Optimal beamforming for model learning

In this section, we present an adaptive method to recover a metal plate geometry without assuming prior knowledge of the propagation model. It is based on focalization in the parameter space (*i.e.* the propagation model is adjusted) and beamforming for localization of the plate boundaries. First, the parametrization of the propagation models using solutions of the Helmholtz equation is presented. Next, we introduce and apply a simple high-pass filter to the beamforming maps to limit the detrimental effect of interference and high-order reflections. A loss function is then designed to assess the focusing capability of a candidate beamformer which should maximize spatial coherency (*i.e.* the energy that is focused at the geometry estimate on the beamforming map) in the case when the propagation model is appropriate. Finally, an optimizer based on simulated annealing is presented to recover optimal propagation parameters by minimizing the loss in a limited number of iterations. The efficiency of such an approach has been successfully demonstrated for underwater source localization purposes [19, 57].

Based on the hypothesis of linear and isotropic propagation, we aim at recovering both a propagation model $\hat{g}(r, t)$ and the plate geometry expressed as a set of lines $\hat{\mathbf{M}} = \{\hat{r}_l, \hat{\alpha}_l\}_{l=1..4}$. For candidate propagation models, we keep relying on the (approximate) solutions of the Helmholtz equation. Hence, estimating the propagation model reduces to the estimation of only three parameters which are the longitudinal velocity c_L , the transverse velocity c_T and the plate thickness d , and that we will gather in the variable $\Omega = \{c_L, c_T, d\}$. This choice is convenient as we constrain the propagation models to be physically plausible while limiting the search space

for the propagation model parameters to a low dimension, which will facilitate the optimization process. As the complete state of the structure may be unknown during an inspection operation, the calibration of these wave parameters may be relevant. Indeed, the plate thickness may not be known accurately all over the surface, as the structural integrity is unknown. Besides, the effects of variation of temperature, pressure, moisture may be interpreted as variations of the effective velocities. For more complex sources of disturbances such as the effect of coupling or structure irregularities, we are not seeking to strictly compensate for their effect with such a simple model, yet sufficiently enough to predict the plate boundaries location accurately. Next, given candidate values Ω for the model parameters, the dispersion equation for the A0 mode is numerically solved, and is used to infer the propagation model $\hat{g}_\Omega(r, t)$ with Eq. (7.1).

Compared to plate geometry reconstruction with a known propagation model, inferring the plate geometry and propagation parameters simultaneously requires the determination of a larger number of unknowns. Consequently, an appropriate loss function that ideally prevents irregular cost surfaces with many local minima is needed to facilitate the optimization process.

Here, we introduce a loss function to assess the ability of a candidate beamformer (related to candidate parameter values Ω) to focus the energy of the wave packets contained in the measurements at the plate geometry estimate, so that minimizing the loss function w.r.t. the model parameters improves the spatial coherency achieved with the beamformer. Given candidate parameters Ω , the beamforming map $\mathcal{L}_\Omega(r, \alpha)$ is constructed as in Eq. (7.2), allowing a plate geometry estimate $\hat{\mathbf{M}}_\Omega$ to be retrieved from it with the same optimization process as in Eq. (7.3). We then evaluate the total energy focused at the estimate $\hat{\mathbf{M}}_\Omega$ over the beamforming map that has been high-pass filtered to limit the effect of interference and high-order reflections. The filtered map value at each line (r, α) is simply defined by the difference between its initial energy value and the minimum of energy in its vicinity that is to be appropriately defined:

$$\tilde{\mathcal{L}}_\Omega(r, \alpha) = \mathcal{L}_\Omega(r, \alpha) - \min_{(r', \alpha') \in V(r, \alpha)} \mathcal{L}_\Omega(r', \alpha'). \quad (7.4)$$

where $V(r, \alpha)$ refers to the set of lines in the vicinity of (r, α) . Compared to the standard $\mathcal{L}_\Omega(r, \alpha)$, the filtered map $\tilde{\mathcal{L}}_\Omega$, while being inexpensive to compute, has the advantage to present fewer areas where the intensity is high but homogeneous (such a situation occurs due to the combination of high-order reflections and interference). Hence, it is more compatible with Eq. (7.3) for the determination of the location of the edges. Thus, using Eq. (7.4) may appropriately filter "fuzzy" areas where the energy is spread homogeneously (*i.e.* not focused at a single point) due to interference on the beamforming maps, and may isolate correct intensity peaks. This will be illustrated next with experimental data.

To simultaneously recover propagation parameters and the plate geometry, performing a joint search in both the propagation model space and the geometry space would be computationally expensive. Instead, we rely on an optimal beamforming formulation: the loss is only evaluated over candidate model parameters Ω , and for the corresponding geometry estimate $\hat{\mathbf{M}}_\Omega$ retrieved from the high-pass filtered beamforming map $\tilde{\mathcal{L}}_\Omega$ with Eq. (7.3). The loss value $l(\Omega)$ that we will seek to minimize is then defined as the opposite of the sum of the intensity levels evaluated at the retrieved edges:

$$l(\Omega) = - \sum_{(r, \alpha) \in \hat{\mathbf{M}}_\Omega} \tilde{\mathcal{L}}_\Omega(r, \alpha), \quad (7.5)$$

so that the energy focused at the plate geometry estimate can be maximized, while the detrimental effects of interference and high-order reflections can be minimized. The geometry estimate provided by the optimal beamformer is expected to match closely the ground truth geometry

due to the maximum of spatial coherency, as long as measurements have been acquired on a sufficient portion of the plate surface. Besides, to highlight the benefit of using high-pass filtered beamforming maps, propagation model selection based on the loss evaluated using the regular beamforming map \mathcal{L}_Ω will also be carried out in the following sections, and the mapping results will be compared.

7.2.3 Optimization with simulated annealing

An optimization process is needed to adjust the wave propagation parameters (longitudinal and transverse velocities and the plate thickness) through the minimization of the loss function. Due to the implicit definition of the propagation models, first-order methods such as gradient descent are impractical to use for our problem.

Simulated annealing is efficient for estimating a global minimum of a cost function, even when the number of unknowns is large [57]. This simple metaheuristic has been successfully used in Chapter 5 for recovering polygonal shapes, and it is a good candidate optimizer for the new problem at hand. For Lamb-wave based mapping without a prior propagation model, the parameter search is reduced to the optimization over the value of Ω , as the geometry is directly deduced using Eq. (7.3) when Ω is fixed. Our implementation of the optimizer is similar to that used in articles related to the underwater localization problem [19]. The major difference lies in the cost function design. First, we initialize the propagation parameters with uniform distributions within predefined intervals:

$$\begin{aligned} c_L^{(0)} &\sim \mathcal{U}([c_{\min}, c_{\max}]); & c_T^{(0)} &\sim \mathcal{U}([c_{\min}, c_L^{(0)} - m]) \\ d^{(0)} &\sim \mathcal{U}([d_{\min}, d_{\max}]). \end{aligned}$$

The value of c_T is drawn below the value of c_L , with a small margin $m > 0$, as it cannot physically be higher (*i.e.* there are no positive real-valued solutions to the Rayleigh-Lamb equation). During each iteration of the algorithm, the parameters are randomly disturbed using the following perturbations:

$$\begin{aligned} c_L^{(p+1)} &= \min \left\{ \max \{ c_L^{(p)} + \Delta c \cdot \chi_0^3, c_{\min} \}, c_{\max} \right\} \\ c_T^{(p+1)} &= \min \left\{ \max \{ c_T^{(p)} + \Delta c \cdot \chi_1^3, c_{\min} \}, c_L^{(p+1)} - m \right\} \\ d^{(p+1)} &= \min \left\{ \max \{ d^{(p)} + \Delta d \cdot \chi_2^3, d_{\min} \}, d_{\max} \right\}. \end{aligned}$$

Using the min-max formulation enables to leverage prior information by constraining the parameter values within plausible intervals. Similarly to the initialization step, $c_T^{(p+1)}$ is restricted to be lower than $c_L^{(p+1)}$ with the same margin m . χ_0, χ_1, χ_2 are independent random values between -1 and 1, and are drawn from uniform distributions: $\chi_{1,2,3} \sim \mathcal{U}([-1, 1])$. They are raised to power 3 so that small variations are more likely. Still, large variations (*i.e.* in the order of Δc for the velocities, and Δd for the thickness) can be possible to explore regions of the search space that are far occasionally from the current $\Omega^{(p)}$.

Next, for the new parameter values $\Omega^{(p+1)}$, the dispersion relation $k_{\Omega^{(p+1)}}(\omega)$ is determined by solving the Raleigh-Lamb equation. As the relation is computed numerically at each iteration of the optimization process, the computational load can be quite demanding. A trick to save computational time is to determine the wavenumber only for a restricted set of pulsation values $\omega_1, \dots, \omega_K$ and to use linear interpolation to determine the wavenumber for other frequencies. The acoustic transfer function $\hat{g}_{\Omega^{(p+1)}}(r, \omega)$ is subsequently determined with the Helmholtz model given in Eq. (7.1). The beamforming map $\mathcal{L}_{\Omega^{(p+1)}}$ is then computed, and a plate geometry

estimate $\hat{\mathbf{M}}_{\Omega^{(p+1)}}$ is retrieved from it, allowing the determination of the loss value $l(\Omega^{(p+1)})$ with Eq. (7.5).

The variation of energy between the previous and disturbed parameters is $\Delta E = l(\Omega^{(p+1)}) - l(\Omega^{(p)})$. In the context of simulated annealing, the disturbed parameter value is not systematically kept. It is the case only when:

$$\chi < \exp \left\{ -\frac{\Delta E}{\gamma T(p)} \right\}$$

where γ is a strictly positive scaling parameter, $T(p)$ is the temperature at the iteration p , and χ is drawn randomly and uniformly between 0 and 1. Consequently, when the energy is decreased, the disturbed parameters are systematically kept. Otherwise, the acceptance rate is given by the Boltzmann distribution, which yields lower acceptance rates for more significant increases of energy. The temperature parameter is often chosen to decrease inverse logarithmically. To enable fast convergence of simulated annealing, we will decrease the temperature inverse linearly. Trials and errors are used to determine appropriate parameters for the optimizer: the values that empirically demonstrate a lower likelihood for the optimizer to be stuck in local maxima while maintaining a sufficient convergence speed are retained.

As it has already been highlighted, the cost function may be heavy to assess, in particular when the number of considered measurements is high. Indeed, each iteration of the optimizer requires solving the Rayleigh-Lamb equation for a set of frequency values, then computing the beamforming map using all the measurements, and recovering a plate geometry estimate with the optimization. For a robotic application, a few seconds might be needed to achieve convergence with around 100 measurements and using a linear decrease of the temperature. Thus, our method cannot be used in real time. Yet, this is not an issue, as it could be considered, during a practical robotic inspection task, to stop the robot for a few seconds to run the optimizer occasionally, and restrict the maximum number of acoustic measurements used to perform the optimization.

7.3 Experimental results

This section illustrates the efficiency of our approach for mapping a rectangular metal panel using ultrasonic guided waves in the three aforementioned scenarios. The benefit of filtering the beamforming maps is first discussed and highlighted based on the experimental data. Next, the correlation between the designed loss function and the reconstruction error is numerically assessed. The optimizer designed in the previous section is run to simultaneously infer the plate geometry and recover the propagation model. The resulting precision is compared with that obtained when using a predetermined propagation model that will be used as a baseline. The overall results demonstrate that this target objective is successfully achieved with our method in nominal acquisition conditions (Scenarios 1 and 2). Our approach is also tested in slightly disturbed conditions (Scenario 3) to illustrate that it remains efficient. For all the situations, the results are obtained using the full batch of measurements, which would amount to having the robot covering an important portion of the plate surface before calibrating the propagation model. As this may not be a representative scenario, the efficiency of our procedure is also assessed with simulated robot paths with acquisition points that are more sparse, and with a varying number of measurements available during the calibration.

7.3.1 Experimental scenarios

We test our approach on experimental acoustic data that have been acquired in three different scenarios. To generate the datasets, an emitter/receiver pair of nearly collocated contact




Figure 7.1: Picture showing the experimental setup for Scenario 3.

piezoelectric transducers is moved to different positions on the plate surface, in each scenario.

For Scenario 1, we use the data acquired manually on our 600x450x6 mm aluminum plate. For Scenario 2, we use the dataset collected on our 1700x1000x6 mm steel plate. For Scenarios 1 and 2, the transducers were in contact with the surface during the acquisition. Coupling gel is placed at their interface to ensure a good coupling. For Scenario 3, the data are acquired on the same plate as the one used for Scenario 1. However, a thin layer of water of approximately 1 mm of thickness is put all over the plate surface, and the transducers are not placed in contact with the plate surface during the acquisition, but are kept in contact with the water only. This acquisition is performed using a customer-design five axes immersion scanner made by Inspection Technology Europe BV. It is also used to place the transducers (that are kept one beside the other) at specific positions on the measurement grid while maintaining a constant distance between the plate surface and the transducers. Due to the plastic holder, the distance between the centers of the transducers is 2.5 cm, which may not be an issue for our approach that relies on an exact pulse-echo setup, as this distance is not large comparing to the signal wavelength. For this scenario, a total of 108 measurements is collected at the same positions as those used for Scenario 1. The sampling frequency is 1.25 MHz, and the number of samples per measurement is 500. The data acquired in this setup are expected to be representative of those that would be acquired on a real structure by the robotic system, where water may also be used as coupling, and where the transducers may not be directly in contact with the structure surface to avoid damaging them by surface irregularities. Furthermore, this scenario will also allow us to highlight the potential of our approach in a slightly "disturbed scenario" comparing to the nominal acquisition conditions of Scenarios 1 and 2. A picture of the experimental setup used for Scenario 3 is available in Fig. 7.1.

7.3.2 Influence of the high-pass filter

First, we illustrate the benefit of applying a high-pass filter to the beamforming maps for more accurate evaluation of the focusing ability of a beamformer. We apply DAS beamforming on all the measurements acquired on the larger steel plate (Scenario 2), and retrieve a map estimate using Eq. (7.3). The propagation model parameters $d = 6$ mm, $c_L = 5880$ m/s and $c_T = 3250$ m/s are used based on prior knowledge of the plate material. The DAS beamforming results are provided in Fig. 7.2a, which depicts two visible intensity peaks. The two other peaks are much




Figure 7.2: Different beamforming maps computed using the data from Scenario 2 and the retrieved edges. a) shows the standard beamforming map, which yields correct estimates. b) depicts the map obtained with the same data sub-sampled in time. One edge is not correctly estimated. c) shows the high-pass filtered beamforming map obtained from the sub-sampled data. The correct plate geometry is recovered, while the fuzzy areas due to interference and higher-order reflections have been partially filtered out.

less visible due to their lower intensity, and can be easily mistaken with interference and high-order reflections, which cause areas of homogeneous intensity. Although the estimation is correct in this case (the estimated plate dimension is 1720x980mm), it may not be robust. To illustrate this, we sub-sample in time the data by a rate of 4 and reconstruct the beamforming map using the same propagation model. The results can be seen in Fig. 7.2b. One may observe that the resulting beamforming map is very similar. However, all but one line boundary are still correctly recovered, causing the plate dimension to be underestimated. This could be detrimental to our final objective, as the resulting map estimate, obtained by maximization of "acoustic focusing" as we defined it, is incorrect.

Using high-pass filtered beamforming maps may increase the robustness of the geometry estimates, and thus, may result in better definition of the focusing ability of DAS beamforming for model calibration. The previous beamforming, after filtering, is shown in Fig. 7.2c. To set up the high pass filter, $V(r, \alpha)$ is chosen as the rectangular area centered at (r, α) with a size $6.5\text{cm} \times 24^\circ$. This window is sufficiently large to encompass the intensity peaks on the beamforming map, and sufficiently limited to make the filtering effective. When now relying on the filtered version of the beamforming map, the map estimation is made correct, even though the detrimental effect of interference and higher-order reflections cannot be completely alleviated. Hence, the introduction of the high-pass filter may provide additional robustness, which is to be leveraged to calibrate the propagation model. It may be even more useful in the presence of disturbances during the acquisition, as will be illustrated in the following.

7.3.3 Correlation between the loss and the mapping accuracy

We assess whether the designed loss is sufficiently correlated with the reconstruction error to ensure that its minimization yields an accurate plate geometry estimate. We create a family of 40 propagation models based on the approximate solutions to the Helmholtz equation, with values of $\Omega = \{c_T, c_L, d\}$ uniformly distributed within the intervals $c_T, c_L \in [2500\text{m/s}, 6500\text{m/s}]$ and $d \in [3\text{mm}, 7\text{mm}]$. Then, we determine the filtered beamforming map using each model to obtain a plate geometry estimate in each of the three scenarios under consideration. The results are displayed in Fig. 7.3 as plots of the reconstruction error (both in range and orientation)

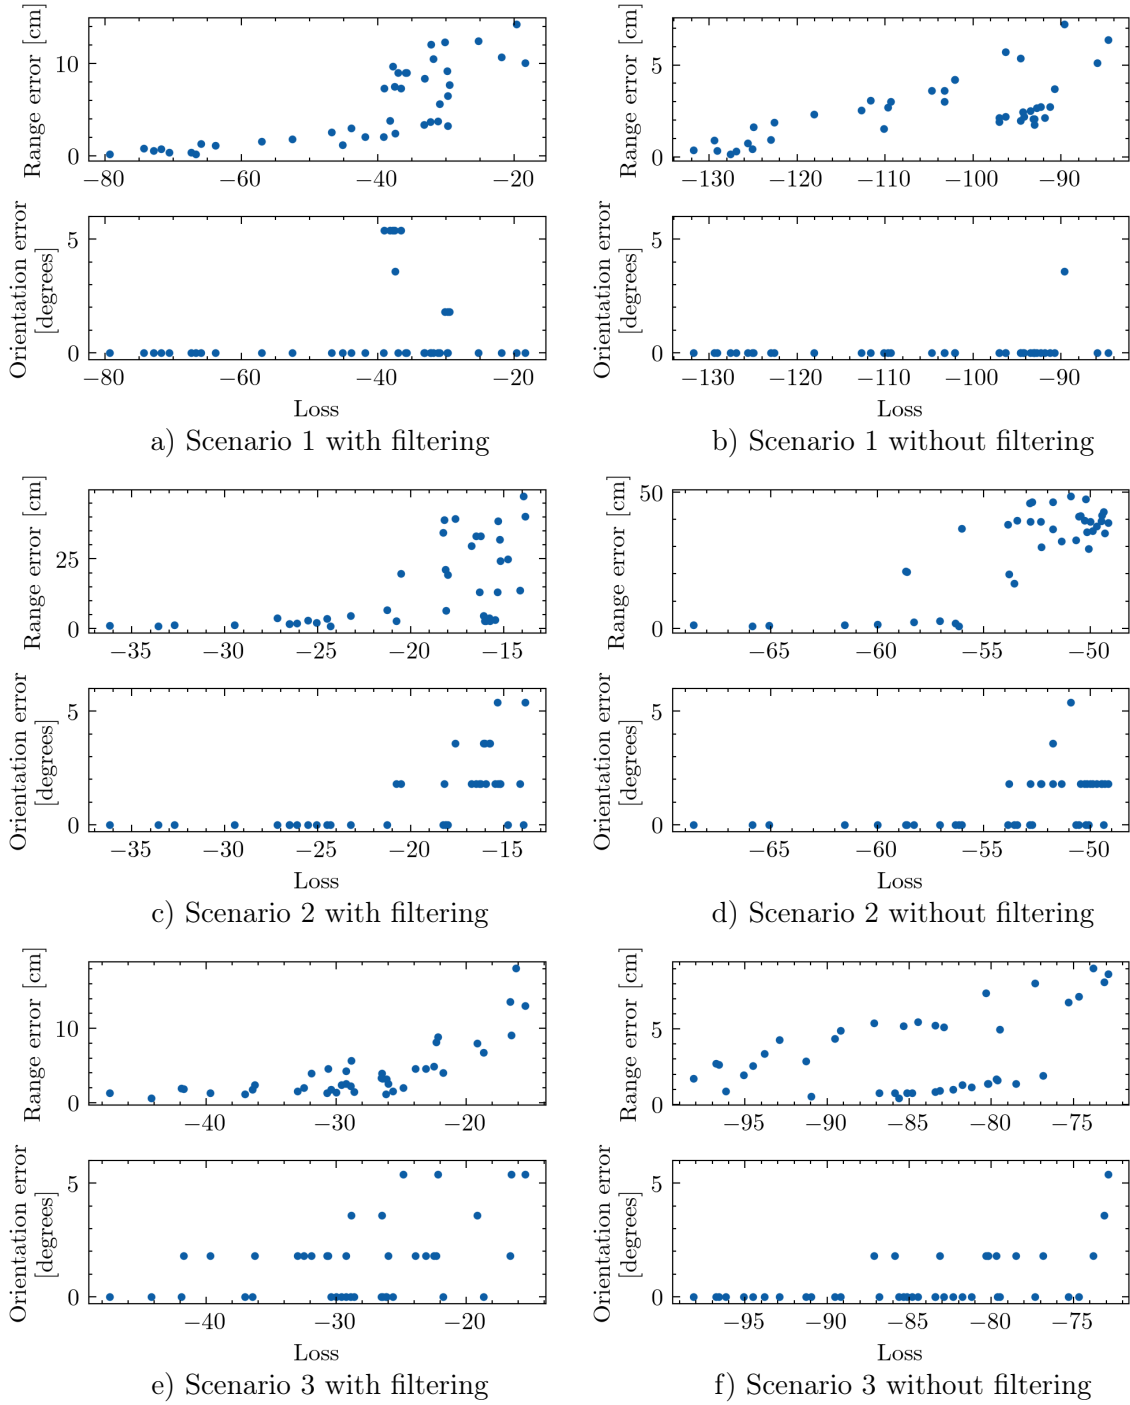


Figure 7.3: 2D plots showing the correlation and loss values achieved using each of the propagation models from the predefined model set, and for the three scenarios. On the left, the loss value and geometry estimates are determined from the high-pass filtered beamforming maps. On the right, they are determined from the standard maps without filtering.

w.r.t. the loss value. We also display the same plots using the geometry estimates and loss values obtained without the filtering, for comparison. Ideally, we would obtain reconstruction errors that would monotonically decrease with the loss value.

For both Scenarios 1 and 2 in Fig. 7.3a-d, we observe that the range errors globally diminish for lower loss values. It can also be observed that models with close loss values can yield completely different reconstruction errors, or that models with high loss values can provide an accurate estimate (*e.g* the plot of the range error in Fig. 7.3c). Yet, these tendencies are not critical, as the major requirement for our approach to be effective is that the best estimation results are achieved for the lower losses as this is what will be minimized. The variations of the orientation error present a different aspect. There is no strict decrease in the error for lower losses. Also, the discretization of the orientation value is visible due to the resolution of the beamforming map that is limited for the sake of computational efficiency. Yet, it is still possible to achieve an accuracy in the order of one degree, which is sufficient for our application. Also, the limited resolution is not a significant problem for propagation model learning, as the best accuracies are also achieved for the lower loss values.

We notice a slightly more divergent behavior for the two losses in Scenario 3, which corresponds to our slightly disturbed scenario. Regarding Fig. 7.3f, where the beamforming maps are not filtered, the reconstruction error presents a tendency to decrease for lower loss values. However, the minimum reconstruction error is achieved for a loss value of approximately $l(\Omega_{\min}) = -86$, and it becomes higher for lower losses. This is not the desired behavior, as it would result in poorer estimation results if the loss minimization is effective. This tendency is not visible when the loss values and the geometry estimates are obtained from the high-pass filtered beamforming maps (Fig. 7.3e), where the lowest reconstruction errors are globally obtained for the lower losses. Hence, these results support the benefit of high-pass filtering the DAS beamforming maps for mapping and propagation model learning. It provides additional robustness due to its ability to lower the detrimental effect of high-order reflection and interference.

7.3.4 Mapping results with model learning in nominal conditions

To assess the performance of our approach for optimal beamforming for Lamb wave-based mapping, we run it 10 times for Scenario 1 and 2, with random initial parameter values for each repetition. The simulated annealing is set with the following parameters: $\Delta c = 7000\text{m/s}$, $\Delta d = 2\text{cm}$, $c_{\min} = 1500\text{m/s}$, $c_{\max} = 8000\text{m/s}$, $m = 500\text{m/s}$, $\gamma = 1$, $T = T_0/t$ where $T_0 = 200$, $d_{\min} = 3\text{mm}$ and $d_{\max} = 1\text{cm}$. Besides, the performance of our approach is compared with the results achieved when using predetermined propagation models based on prior knowledge of the plate thickness, the longitudinal and transversal velocities in nominal conditions. These predetermined models will serve as a baseline for comparison. Fig. 7.4 depicts the evolution of the minimum loss value and the corresponding reconstruction errors achieved during the optimization process. The upper and lower bounds of the filled areas represent the minimum and maximum values over the 10 runs, at each iteration step, while the solid blue lines account for the mean values. For comparison, the loss values and reconstruction errors achieved when using the predetermined propagation models for the two scenarios are also displayed.

In both scenarios, the achieved minimum loss decreases monotonically until it reaches a plateau. The final loss value is always lower than the loss value achieved with the predetermined model in Scenario 1 and significantly lower in Scenario 2. In terms of reconstruction error, the final average range error is comparable to that achieved with the predetermined model, which is 3mm for Scenario 1. It is slightly above the estimation error value obtained with the predetermined model (0.5cm) for Scenario 2. Besides, the orientation errors rapidly decrease

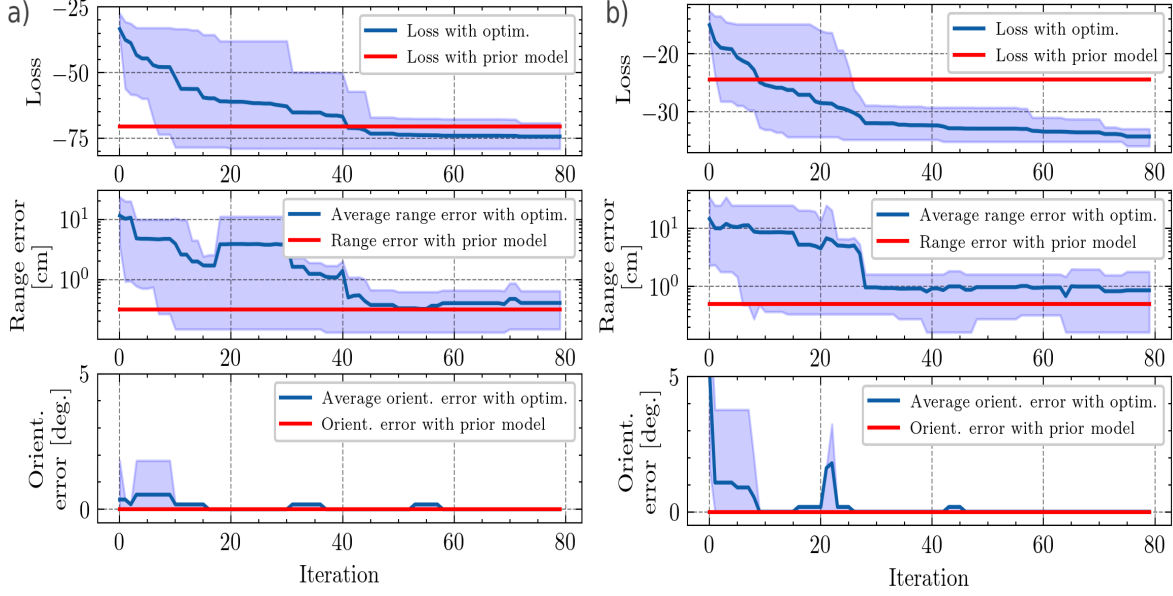


Figure 7.4: Evolution of the loss value and reconstruction errors for 10 repetitions of the optimization process with simulated annealing, and using the data from Scenario 1 (a) and Scenario 2 (b). The solid lines represent the mean values. The upper and lower bounds of the filled areas represent the minimum and maximum values respectively during each iteration. For comparison, the values achieved with the predetermined propagation model for each scenario are displayed as horizontal lines. The scale along the y-axis is logarithmic for the range error plots only.

to zero for all the runs, achieving the same precision as the predetermined models in the two scenarios. Altogether, these experiments demonstrate that a propagation model can be efficiently recovered to estimate a plate geometry through optimal beamforming by relying directly on data, regardless of the plate size and material, and as long as the measurements have been acquired on a sufficient portion of the surface.

7.3.5 Mapping results with model learning in non-nominal conditions

We perform the same experiments using the data from Scenario 3 which have been acquired in slightly disturbed conditions due to the layer of water that is placed at the interface between the transducers and the plate surface. The evolution of the loss and the reconstruction errors for 10 repetitions of the optimization process with random initial parameter values are shown in Fig. 7.5. It can be observed that the loss rapidly decreases to a value lower than that achieved with the predetermined model. Regarding the reconstruction error, the average range error, after convergence, reaches 1cm while the error is 2cm when using the predetermined model for the geometry reconstruction. Furthermore, the orientation error rapidly decreases to a value lower than one degree.

Altogether, our approach, which automatically calibrates the propagation model based on data, yields better geometry estimates in this disturbed scenario compared to mapping with the predetermined propagation model, as illustrated by Fig. 7.5. These results are promising as the propagation in the thin layer of water has not been modeled explicitly. This further shows the potential of our approach which adapts the propagation model to achieve accurate plate reconstruction in varying conditions, as it is expected to happen for a robotic inspection

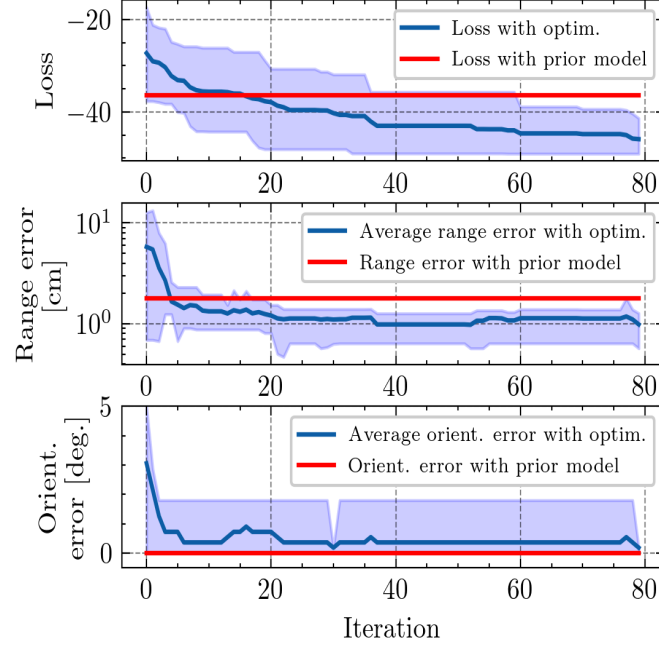


Figure 7.5: Evolution of the loss value and reconstruction errors for 10 repetitions of the optimization process and using the data from Scenario 3. The blue lines represent the mean values. The upper and lower bounds of the colored areas represent the minimum and maximum values respectively during each iteration. For comparison, the values achieved with the predetermined propagation model are displayed as horizontal red lines. The scale along the y-axis is logarithmic for the range error plot only.

task on large structures in challenging outdoor environments. Also, with our approach, there is potentially no need to take into account slight sources of disturbance explicitly in the propagation models.

7.3.6 Evaluation of our approach with sparse measurements

During a real robotic inspection task, the acoustic measurements may not have been acquired over a dense grid, and over a sufficient portion of the surface during the calibration operation to fully recover the plate geometry. To evaluate the performance of our approach in more realistic scenarios, we assess how the quantity of measurements available during the optimization process affects the propagation model selection when the data are acquired along more realistic trajectories. To avoid running the optimizer every time, we consider the previous family of 40 propagation models. We evaluate, for each propagation model, the reconstruction errors obtained with and without filtering the beamforming maps, and select the error relative to the propagation model, for every set of measurements, that yields the lowest loss value. We also assess the estimation errors when the predetermined propagation model is used along with high-pass filtered beamforming maps, for comparison. We design two different trajectories (one for each plate) that are provided in Fig. 7.6. The results, which are shown in Fig. 7.7, were obtained using the data acquired on the large steel plate (Scenario 2) and using the data acquired in the disturbed conditions (Scenario 3).

When only a few measurements are available, the average range errors are relatively high in




Figure 7.6: Simulated robot paths used for the experiments on the steel plate (left) and the aluminum plate (right). The path lengths (*i.e* the number of acquisition positions) are respectively 36 and 26.

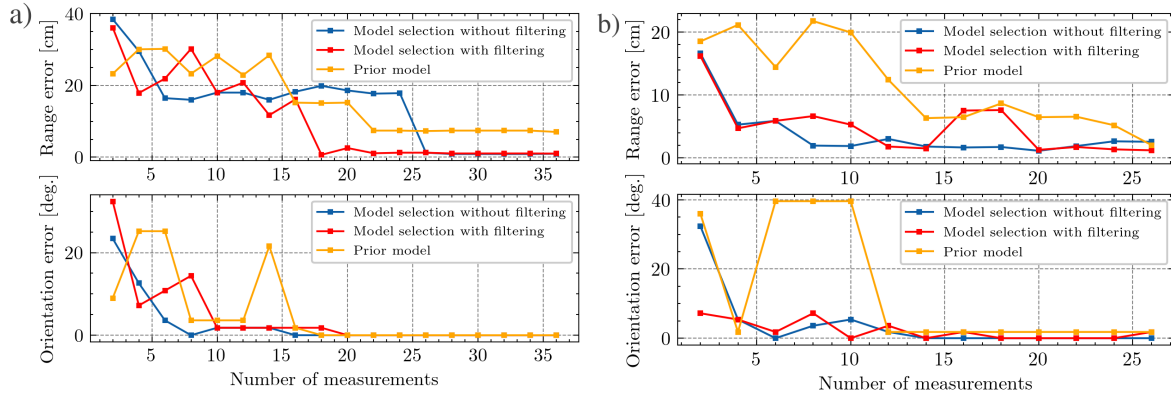


Figure 7.7: Reconstruction errors w.r.t. the number of measurements considered in the loss, and using the data from Scenario 2 (a) or from Scenario 3 (b). The measurements are integrated one after the other in the loss, following the simulated paths. The error is only evaluated for the model yielding the lower loss value for the two cases (model selection with and without filtering the beamforming maps). For comparison, the errors achieved with the predetermined models are also displayed.




Figure 7.8: Loss value computed on the high-pass filtered beamforming maps with respect to the longitudinal and transverse velocities. The data from Scenario 1 are used for the plot on the left, while the data from Scenario 2 are used for the plot on the right.

all the cases, because the sensors need to pass by a border closely enough to detect it. The most accurate results are achieved when all the measurements are taken into account. In the middle, we can see, in Fig. 7.7a, that the minimum range error is reached faster when the beamforming maps have been filtered, and this minimum error is lower than the error achieved with the predetermined model. Furthermore, both in Fig. 7.7a and Fig. 7.7b, the reconstruction errors are seemingly lower when relying on model calibration for the reconstruction. Overall, the results show that our approach is effective for mapping, and yields similar - if not better - performance than that obtained with a predetermined propagation model, even when the measurements have been sparsely acquired by the mobile unit.

7.3.7 Discussions

Although we have shown that a propagation model can be recovered based on data to achieve precise localization by adapting the model parameters, our method cannot recover the real values of the physical model parameters. Indeed, different model parameters can yield equivalent loss values as illustrated by Fig. 7.8, where we represent the variation of the loss w.r.t. the two velocity parameters for a fixed thickness ($d = 6\text{mm}$), and using the data acquired on the two metal plates considered in this study (Scenario 1 and 2). The more specific explanation is that different model parameters may lead to similar dispersion values in the considered frequency bandwidth, after resolving the Rayleigh-Lamb equations. Yet, the fact that the actual model parameters cannot be recovered with our method is not a significant issue, as our primary objective is to achieve accurate mapping of the plate structure with a calibrated model g , so that the wavepackets inside the measurements can be appropriately accounted for.

The method presented here restricts the propagation models to be approximate solutions of the Helmholtz equation with dispersion relations derived from the Rayleigh-Lamb equation. This approach is not sufficient for recovering the actual model parameters, and addressing this in future work would be desirable for a complete NDT method (the knowledge of plate thickness variation would be of particular importance). Yet, the recovered propagation model \hat{g} is sufficient to achieve accurate mapping results through beamforming. Whether this approach remains appropriate in more realistic conditions (*i.e* on a real ship hull for example) is still subject to investigation. It is expected that this parametrization would be sufficient as long as the

first-order wave-packets are sufficiently energetic within the measurements, and that the major hypotheses on the propagation model (linear propagation, homogeneous and isotropic material) approximately hold. Due to the large size of the plates on a real structure and the low reflectivity of weld joints, the signal-to-noise ratio (SNR) is expected to be low. Hence, a study evaluating the performance of our approach for various SNR conditions is needed to assess how likely it would work in practice. Also, it is to be noted that, in the present study, the sensor positions were measured accurately, whereas in practice, only estimated positions will be available.

The case of multi-modal propagation is not considered here, whereas it is likely if the frequency is not sufficiently adequate to the material in a -presumably- unknown state, or if mode conversion occurs. Having at least A0 and S0 modes propagating simultaneously is the most frequent scenario. It is believed that the algorithm could be extended by considering several hypotheses (bi-modal, A0-only, S0-only propagation, mode conversion for different paths, see [108]) to determine which one is most likely based on data. However, integrating more complex interactions such as diffraction due, for example, to complex structural features such as stiffeners, holes... would be more challenging. We expect isotropic propagation to be prevalent, as these complex wave interactions may be scarce and have a sufficiently small incidence on the signals to not affect the mapping results.

Besides, our method, as presented here, is restricted to rectangular geometries, as it facilitates edge retrieval from the beamforming map. This constraint is not a limitation for mapping storage tanks or ship hulls, as they are almost entirely made of rectangular panels. Yet, our approach could be extended, following the framework presented in Chapter 5, to make it applicable to structures with more general geometries. One may also want to adapt our approach to more conventional applications in SHM where, for example, the propagation model could be automatically calibrated by maximizing the energy focused at the estimated defect location on the imaging results, or by relying on the reflections on the sample boundaries.

7.4 Chapter summary

This chapter introduces a method to accurately recover the geometry of a metal plate by relying on ultrasonic measurements acquired by a mobile unit, in a pulse-echo mode, and without using a predetermined propagation model. Our approach is based on focalization in the model parameter space and DAS beamforming for localization of the plate boundaries. We restrict the propagation model to be an approximate solution of the Helmholtz equation and parameterize it with only three physical values. We introduce a loss function assessed on high-pass filtered beamforming maps to quantify the focusing ability of a candidate beamformer (*i.e.*, candidate model parameters). Eventually, we find optimal model parameters with a simulated annealing optimization process.

We demonstrate the performance of our method on three sets of experimental acoustic data acquired on a dense grid on two metal plates of different sizes and different materials. The results illustrate the relevance of filtering the beamforming maps to reduce the detrimental effect of interference and high-order reflections in the assessment of the focusing ability of beamforming. We also show that a propagation model enabling accurate boundary localization can be recovered with simulated annealing. For the two undisturbed scenarios, the precision of the localization is found to be similar to that achieved with the model built from prior knowledge on the plate material, but it is found to be superior in non-nominal acquisition conditions. This highlights the very potential of our method for Lamb wave-based localization and mapping on a large metal structure, where the wave propagation conditions may not be known a priori. The

benefit of this approach is real for practical industrial inspection tasks, where the propagation models could be automatically calibrated. Eventually, we assess the performance of our approach using reduced numbers of measurements acquired on the two plates, and following more realistic robot trajectories. The results illustrate that the proposed approach can recover the plate geometry accurately even with sparse measurements, and outperforms the mapping based on the predetermined propagation model.

Yet, the method shall be integrated within a simultaneous localization a mapping framework, as the sensor positions need to be estimated as well. Also, the hypothesis of rectangular plate geometries shall be relaxed, and more complex wave phenomena such as anisotropic and/or multi-modal propagation, diffraction, or mode conversion are to be integrated in the model. The recovery of the real physical parameters, such as plate thickness, and the mapping of defects shall be investigated to make possible a complete robotic NDT task. Finally, the method shall be tested in more realistic conditions.

Chapter 8

Conclusion

8.1 Summary of the contributions

This thesis elaborates various algorithmic frameworks to evaluate the potential of omnidirectional ultrasonic guided waves for localization and mapping purposes on large metal structures, such as storage tanks or ship hulls. The underlying motivation is to contribute to the emergence of (potentially autonomous) mobile robotic systems for accurately mapping large surfaces: by processing the long-range ultrasonic measurements acquired over the robot trajectory, the presence of potential defects that may act as acoustic scatterers could eventually be revealed.

As the knowledge of the measurement positions is critical for such an application, we first address the localization-only problem on a rectangular metal panel by relying on the ultrasonic reflections of an excitation wave on the plate boundaries. We propose a simple particle filter algorithm whose observation model is obtained by correlating the ultrasonic measurements with a wave propagation model. The particularity of this approach is that, on the contrary to the mapping strategies reviewed in the literature, no explicit echo detection nor echo association with the plate boundaries are required, whereas these are challenging tasks in signal processing, especially in the presence of noise. The experimental results obtained in a laboratory environment demonstrate the potential of this approach for accurate positioning.

We subsequently address the corresponding SLAM problem to additionally recover the position of the plate boundaries. The proposed method is a FastSLAM algorithm that integrates Delay-and-Sum (DAS) beamforming to solve the online SLAM problem. The performance of this approach is compared with that of a second FastSLAM algorithm that relies on echo detection and echo-plate boundary association. The experiments demonstrate that the former method provides more robust localization and mapping results. A SLAM experiment with a magnetic crawler is also carried out to demonstrate the applicability of ultrasonic SLAM with a real system.

Despite its apparent simplicity, DAS beamforming, which can be used for mapping acoustic scatterers such as the boundaries of a metal panel, presents several limitations, namely signal interference, and the lack of uncertainty assessment on the map estimates. To address these issues, we develop a more elaborate mapping framework that can jointly recover a probabilistic grid representation along with a feature-based representation of the map. In this context, the consistency between the two representations can be leveraged to mitigate the effect of interference. Furthermore, this novel mapping framework is integrated within a frontier exploration strategy for autonomous map reconstruction, and it is also leveraged for recovering more arbitrary polygonal shapes than rectangles. The performance of the developed methods is extensively evaluated in simulation, and experimentally validated using real robotic platforms.

In the next chapter, we propose a GraphSLAM framework to solve the full SLAM problem. In our case-study, it consists in simultaneously recovering the metal plate geometry and the sensor trajectory through batch optimization, using all the available ultrasonic measurements. We first propose a novel observation model that can leverage the high-order reflections for localization. The latter model is integrated, along with DAS beamforming, into a NLS optimizer. The experiments illustrate the benefits of the proposed approach, as it can achieve more accurate SLAM results than FastSLAM in more challenging situations.

Finally, as the SLAM methods we have elaborated essentially rely on an accurate wave propagation model, the last chapter of this thesis presents an optimal beamforming approach for mapping a metal plate with little prior knowledge on the model, as it may be lacking during a practical inspection task. The experiments performed in different conditions show that, not only the proposed approach is effective, but it can lead to better mapping results, in non-nominal acquisition conditions, when compared to mapping with a predetermined propagation model. Indeed, the latter may no longer yield accurate results in the presence of variations of the acquisition conditions, for example.

8.2 Perspectives

Despite the promising results on ultrasonic SLAM that we were capable to obtain, it is clear that the development of a true long-range inspection robot is a challenging task which necessitates substantial work to emerge.

Although our work proposes algorithmic frameworks for mapping acoustic scatterers with a mobile robotic system, one may note that no experiments are performed with the aim to detect actual defects. The reason is that the detection of a defect signature within a multi-echo ultrasonic measurement is challenging. This is all the more true due to multi-modal propagation of the waves, and the complex scattering effects that may occur in structures with complex geometries. Hence, future works shall focus on how the signals related to geometrical structures, those related to defects, and the signals due to noise can efficiently be identified and separated for subsequent defect identification. A promising research direction in that regard would be to explore the use of more advanced and adaptive beamforming methods, instead of the simple delay-and-sum approach, such as Minimum Variance Distorsionless Response (MVDR) beamforming, or subspace-based methods.

Another remark on the present work is that different SLAM problems are addressed in this thesis, but it is not specified exactly how they should be merged together. For example, it is not yet explicit how our FastSLAM and GraphSLAM frameworks can leverage optimal beamforming to calibrate the propagation model, and whether the mapping framework with combined representations can be easily integrated as well. The actual reason is that the different methods need first to be evaluated in realistic conditions, on actual large structures, so that one may better understand their limitations for the target application. Finally, very little work in this thesis focuses on instrumentation, whereas the efficient generation and acquisition of the signals is a major stake for long-range inspection. This also ought to be further investigated.

Appendix: Derivation of the integrals for calculation of the inverse measurement models

In this appendix, we provide calculation details on how the integrals needed to determine inverse models in Section 5.2 are determined.

A: Derivation of the first integral

To derive the expression of the first integral in Eq. 5.2.4, we need to show that, for $0 \leq r \leq a$:

$$\int h(r - \rho) \mathcal{N}_{[0,a]}(\rho | \mu, \sigma^2) d\rho = \frac{\Phi_{\mu,\sigma}(r) - \Phi_{\mu,\sigma}(0)}{\Phi_{\mu,\sigma}(a) - \Phi_{\mu,\sigma}(0)},$$

where $\Phi_{\mu,\sigma}$ is expressed with the error function:

$$\Phi_{\mu,\sigma}(\rho) = \text{erf}\left(\frac{\rho - \mu}{\sigma\sqrt{2}}\right).$$

First, we make explicit the normalization factor of the truncated Gaussian distribution:

$$\eta = \int_0^a \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left\{-\frac{(\rho - \mu)^2}{2\sigma^2}\right\} d\rho = F(a) - F(0),$$

where F is the cumulative distribution of a Gaussian distribution with mean μ and standard deviation σ : $F(x) = \frac{1}{2} \left[1 + \text{erf}\left(\frac{x - \mu}{\sigma\sqrt{2}}\right)\right] = \frac{1}{2} [1 + \Phi_{\mu,\sigma}(x)]$. Finally, we have:

$$\int h(r - \rho) \mathcal{N}_{[0,a]}(\rho | \mu, \sigma^2) d\rho = \frac{1}{\eta} \int_0^r \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left\{-\frac{(\rho - \mu)^2}{2\sigma^2}\right\} d\rho = \frac{F(r) - F(0)}{F(a) - F(0)},$$

that yields, when substituting the expression of F , the desired result.

B: Derivation of the second integral

To derive the expression of the second integral in Eq. 5.2.4, we need to show that:

$$\int h(r - \rho) \mathcal{N}_{[0,a]}(\rho | \mu, \sigma^2) \exp\left\{-\frac{(\rho - \nu)^2}{2b^2}\right\} d\rho = \frac{\Gamma_{\mu_{\text{new}},\sigma_{\text{new}}}(r) - \Gamma_{\mu_{\text{new}},\sigma_{\text{new}}}(0)}{\Phi_{\mu,\sigma}(a) - \Phi_{\mu,\sigma}(0)},$$

where we have:

$$\begin{aligned} \Gamma_{\mu_{\text{new}},\sigma_{\text{new}}}(x) &= \frac{\sigma_{\text{new}}}{\sigma} \cdot \exp\left(-\frac{(\mu_{\text{new}} - \mu)^2}{2\sigma^2} - \frac{(\mu_{\text{new}} - \nu)^2}{2b^2}\right) \cdot \Phi_{\mu_{\text{new}},\sigma_{\text{new}}}(x), \\ \mu_{\text{new}} &= \frac{b^2\mu + \sigma^2\nu}{b^2 + \sigma^2}, \quad \sigma_{\text{new}} = \frac{\sigma b}{\sqrt{b^2 + \sigma^2}}. \end{aligned}$$

Given the fact that the product of two Gaussian probability density functions is still Gaussian, and based on the expressions of μ_{new} and σ_{new} , we have:

$$\exp\left\{-\frac{(\rho - \mu)^2}{2\sigma^2} - \frac{(\rho - \nu)^2}{2b^2}\right\} = \lambda \cdot \exp\left\{-\frac{(\rho - \mu_{\text{new}})^2}{2\sigma_{\text{new}}^2}\right\},$$

where λ is determined by assessing the two sides of the equation at $\rho = \mu_{\text{new}}$ which yields:

$$\lambda = \exp \left(-\frac{(\mu_{\text{new}} - \mu)^2}{2\sigma^2} - \frac{(\mu_{\text{new}} - \nu)^2}{2b^2} \right)$$

. Hence:

$$\begin{aligned} & \int h(r - \rho) \mathcal{N}_{[0,a]}(\rho|\mu, \sigma^2) \exp \left\{ -\frac{(\rho - \nu)^2}{2b^2} \right\} d\rho \\ &= \frac{\lambda}{\eta} \int_0^r \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left\{ -\frac{(\rho - \mu_{\text{new}})^2}{2\sigma_{\text{new}}^2} \right\} d\rho \\ &= \frac{\lambda\sqrt{2\pi\sigma_{\text{new}}^2}}{\eta\sqrt{2\pi\sigma^2}} \int_0^r \frac{1}{\sqrt{2\pi\sigma_{\text{new}}^2}} \exp \left\{ -\frac{(\rho - \mu_{\text{new}})^2}{2\sigma_{\text{new}}^2} \right\} d\rho \\ &= \lambda \frac{\sigma_{\text{new}}}{\sigma} \frac{\Phi_{\mu_{\text{new}}, \sigma_{\text{new}}}(r) - \Phi_{\mu_{\text{new}}, \sigma_{\text{new}}}(0)}{\Phi_{\mu, \sigma}(a) - \Phi_{\mu, \sigma}(0)} \end{aligned}$$

which yields the desired result after substituting the products $\lambda \cdot \sigma_{\text{new}} \cdot \Phi_{\mu_{\text{new}}, \sigma_{\text{new}}} / \sigma$ with $\Gamma_{\mu_{\text{new}}, \sigma_{\text{new}}}$.

Bibliography

- [1] Bugwright2, autonomous robotic inspection and maintenance on ship hulls and storage tanks, description of the innovative action, 2019.
- [2] Jont B Allen and David A Berkley. Image method for efficiently simulating small-room acoustics. *The Journal of the Acoustical Society of America*, 65(4):943–950, 1979.
- [3] David N Alleyne and Peter Cawley. The interaction of lamb waves with defects. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 39(3):381–397, 1992.
- [4] Francesco Amigoni and Alessandro Gallo. A multi-objective exploration strategy for mobile robots. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, pages 3850–3855. IEEE, 2005.
- [5] Franz Andert and Lukas Goormann. Combined grid and feature-based occupancy map building in large outdoor environments. In *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2065–2070. IEEE, 2007.
- [6] Steven R Anton, Daniel J Inman, and Gyuhae Park. Reference-free damage detection using instantaneous baseline measurements. *AIAA journal*, 47(8):1952–1964, 2009.
- [7] Arthur B Baggeroer and William A Kuperman. Matched field processing in ocean acoustics. In *Acoustic signal processing for ocean exploration*, pages 79–114. Springer, 1993.
- [8] Joan Bordoy, Christian Schindelhauer, Fabian Höflinger, and Leonhard M Reindl. Exploiting acoustic echoes for smartphone localization and microphone self-calibration. *IEEE Transactions on Instrumentation and Measurement*, 69(4):1484–1492, 2019.
- [9] Jeffrey Borish. Extension of the image model to arbitrary polyhedra. *The Journal of the Acoustical Society of America*, 75(6):1827–1836, 1984.
- [10] Frederic Bourgault, Alexei A Makarenko, Stefan B Williams, Ben Grocholsky, and Hugh F Durrant-Whyte. Information based adaptive robotic exploration. In *IEEE/RSJ international conference on intelligent robots and systems*, volume 1, pages 540–545. IEEE, 2002.
- [11] Peter Cawley and David Alleyne. The use of lamb waves for the long range inspection of large structures. *Ultrasonics*, 34(2-5):287–290, 1996.
- [12] Sonda Chaabene, Faker Bouchoucha, Mohamed Najib Ichchou, and Mohamed Haddar. Wave mode diffusion and propagation in structural wave guide under varying temperature. *Applied Acoustics*, 108:84–91, 2016.
- [13] CY Chang and FG Yuan. Extraction of guided wave dispersion curve in isotropic and anisotropic materials by matrix pencil method. *Ultrasonics*, 89:143–154, 2018.

- [14] Qi Chen, Kailiang Xu, and Dean Ta. High-resolution lamb waves dispersion curves estimation and elastic property inversion. *Ultrasonics*, 115:106427, 2021.
- [15] Zhe Chen et al. Bayesian filtering: From kalman filters to particle filters, and beyond. *Statistics*, 182(1):1–69, 2003.
- [16] Mei Yi Cheung, Dehann Fourie, Nicholas R Rypkema, Pedro Vaz Teixeira, Henrik Schmidt, and John Leonard. Non-gaussian slam utilizing synthetic aperture sonar. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3457–3463. IEEE, 2019.
- [17] Jinwoo Choi, Yeongjun Lee, Taejin Kim, Jongdae Jung, and Hyun-Taek Choi. Ekf slam using acoustic sources for autonomous underwater vehicle equipped with two hydrophones. In *OCEANS 2016 MTS/IEEE Monterey*, pages 1–4. IEEE, 2016.
- [18] See Yenn Chong and Michael D Todd. Dispersion curve estimation via a spatial covariance method with ultrasonic wavefield imaging. *Ultrasonics*, 89:46–63, 2018.
- [19] Michael D Collins and WA Kuperman. Focalization: Environmental focusing and source localization. *The Journal of the Acoustical Society of America*, 90(3):1410–1422, 1991.
- [20] Alan R Curtis, Michael JD Powell, and John K Reid. On the estimation of sparse jacobian matrices. *J. Inst. Math. Appl*, 13(1):117–120, 1974.
- [21] Frank Dellaert. Factor graphs and gtsam: A hands-on introduction. Technical report, Georgia Institute of Technology, 2012.
- [22] Frank Dellaert, Dieter Fox, Wolfram Burgard, and Sebastian Thrun. Monte carlo localization for mobile robots. In *Proceedings 1999 IEEE international conference on robotics and automation (Cat. No. 99CH36288C)*, volume 2, pages 1322–1328. IEEE, 1999.
- [23] Gordon Dobie, S Gareth Pierce, and Gordon Hayward. The feasibility of synthetic aperture guided wave imaging to a mobile sensor platform. *Ndt & E International*, 58:10–17, 2013.
- [24] Ivan Dokmanić, Laurent Daudet, and Martin Vetterli. From acoustic room reconstruction to slam. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6345–6349. Ieee, 2016.
- [25] Rachel S Edwards, Charles Macleod, and Stephen G Pierce. Guided wave based-occupancy grid robotic mapping. In *European Workshop on Structural Health Monitoring: Special Collection of 2020 Papers-Volume 2*, volume 2, page 267. Springer Nature, 2021.
- [26] Alberto Elfes. Using occupancy grids for mobile robot perception and navigation. *Computer*, 22(6):46–57, 1989.
- [27] Baofu Fang, Jianfeng Ding, and Zaijun Wang. Autonomous robotic exploration based on frontier point optimization and multistep path planning. *IEEE Access*, 7:46104–46113, 2019.
- [28] Eric B Flynn, Michael D Todd, Paul D Wilcox, Bruce W Drinkwater, and Anthony J Croxford. Maximum-likelihood estimation of damage location in guided-wave structural health monitoring. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 467(2133):2575–2596, 2011.

-
- [29] José M Galán and Ramón Abascal. Numerical simulation of lamb wave scattering in semi-infinite plates. *International Journal for Numerical Methods in Engineering*, 53(5):1145–1173, 2002.
- [30] Lal C Godara. Application of antenna arrays to mobile communications. ii. beam-forming and direction-of-arrival considerations. *Proceedings of the IEEE*, 85(8):1195–1245, 1997.
- [31] Mohammad-Hossein Golbon-Haghighi. Beamforming in wireless networks. *InTech Open*, pages 163–192, 2016.
- [32] Cesar Gonzales, Artur Balyan, Marvin Mallari, and Christoph Schaal. Lamb wave-based nondestructive inspection using a mobile robotic platform. In *Health Monitoring of Structural and Biological Systems XV*, volume 11593, page 115932V. International Society for Optics and Photonics, 2021.
- [33] Héctor H González-Banos and Jean-Claude Latombe. Navigation strategies for exploring indoor environments. *The International Journal of Robotics Research*, 21(10-11):829–848, 2002.
- [34] Michael Grant, Stephen Boyd, and Yinyu Ye. Cvx: Matlab software for disciplined convex programming, 2008.
- [35] Giorgio Grisetti, Rainer Kümmerle, Cyrill Stachniss, and Wolfram Burgard. A tutorial on graph-based slam. *IEEE Intelligent Transportation Systems Magazine*, 2(4):31–43, 2010.
- [36] Kexin Guo, Zhirong Qiu, Cunxiao Miao, Abdul Hanif Zaini, Chun-Lin Chen, Wei Meng, and Lihua Xie. Ultra-wideband-based localization for quadcopter navigation. *Unmanned Systems*, 4(01):23–34, 2016.
- [37] James S Hall, Peter McKeon, L Satyanarayan, Jennifer E Michaels, Nico F Declercq, and Yves H Berthelot. Minimum variance guided wave imaging in a quasi-isotropic composite plate. *Smart Materials and Structures*, 20(2):025013, 2011.
- [38] James S Hall and Jennifer E Michaels. Multipath ultrasonic guided wave imaging in complex structures. *Structural Health Monitoring*, 14(4):345–358, 2015.
- [39] Joel B Harley and Chen Ciang Chia. Statistical partial wavefield imaging using lamb wave signals. *Structural Health Monitoring*, 17(4):919–935, 2018.
- [40] John Heading. *Mathematical methods in science and engineering*. Elsevier Publishing Company, 1970.
- [41] Wolfgang Herbordt and Walter Kellermann. Adaptive beamforming for audio signal acquisition. In *Adaptive Signal Processing*, pages 155–194. Springer, 2003.
- [42] Erin Hong and Christoph Schaal. Reverse engineering stiffened plates using guided wave-based nondestructive testing methods. In *Health Monitoring of Structural and Biological Systems XII*, volume 10600, page 106000E. International Society for Optics and Photonics, 2018.
- [43] Peter Huthwaite and Francesco Simonetti. High-resolution guided wave tomography. *Wave Motion*, 50(5):979–993, 2013.

- [44] Don H Johnson and Dan E Dudgeon. *Array signal processing: concepts and techniques*. Simon & Schuster, Inc., 1992.
- [45] Miguel Juliá, A. Gil, and Óscar Reinoso. A comparison of path planning strategies for autonomous exploration and mapping of unknown environments. *Autonomous Robots*, 33:427–444, 2012.
- [46] Michael Kaess, Hordur Johannsson, Richard Roberts, Viorela Ila, John J Leonard, and Frank Dellaert. isam2: Incremental smoothing and mapping using the bayes tree. *The International Journal of Robotics Research*, 31(2):216–235, 2012.
- [47] Michael Kaess, Ananth Ranganathan, and Frank Dellaert. isam: Incremental smoothing and mapping. *IEEE Transactions on Robotics*, 24(6):1365–1378, 2008.
- [48] Carl T Kelley. *Iterative methods for optimization*. SIAM, 1999.
- [49] Scott Kirkpatrick, C Daniel Gelatt Jr, and Mario P Vecchi. Optimization by simulated annealing. *science*, 220(4598):671–680, 1983.
- [50] Miranda Kreković, Gilles Baechler, Ivan Dokmanić, and Martin Vetterli. Structure from sound with incomplete data. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3539–3543. IEEE, 2018.
- [51] Miranda Kreković, Ivan Dokmanić, and Martin Vetterli. Echoslam: Simultaneous localization and mapping with acoustic echoes. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 11–15. Ieee, 2016.
- [52] Miranda Krekovic, Ivan Dokmanic, and Martin Vetterli. Look, no beacons! optimal all-in-one echoslam. *arXiv preprint arXiv:1608.08753*, 2016.
- [53] Miranda Kreković, Ivan Dokmanić, and Martin Vetterli. Omnidirectional bats, point-to-plane distances, and the price of uniqueness. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3261–3265. Ieee, 2017.
- [54] Andrii Kulakovskiy. *Development of a SHM system by elastic guided waves applied to aeronautic structures*. PhD thesis, Université Paris Saclay (COmUE), 2019.
- [55] Rainer Kümmerle, Giorgio Grisetti, Hauke Strasdat, Kurt Konolige, and Wolfram Burgard. g 2 o: A general framework for graph optimization. In *2011 IEEE International Conference on Robotics and Automation*, pages 3607–3613. IEEE, 2011.
- [56] Tribikram Kundu, X Yang, H Nakatani, and N Takeda. A two-step hybrid technique for accurately localizing acoustic source in anisotropic structures without knowing their material properties. *Ultrasonics*, 56:271–278, 2015.
- [57] William A Kuperman, Michael D Collins, John S Perkins, and NR Davis. Optimal time-domain beamforming with simulated annealing including application of apriori information. *The Journal of the Acoustical Society of America*, 88(4):1802–1810, 1990.
- [58] Heinrich Kuttruff. *Room acoustics*. Crc Press, 2016.
- [59] Horace Lamb. On waves in an elastic plate. *Proceedings of the Royal Society of London. Series A, Containing papers of a mathematical and physical character*, 93(648):114–128, 1917.

-
- [60] Sang Jun Lee, Navneet Gandhi, James S Hall, Jennifer E Michaels, Buli Xu, Thomas E Michaels, and Massimo Ruzzene. Baseline-free guided wave imaging via adaptive source removal. *Structural Health Monitoring*, 11(4):472–481, 2012.
- [61] Simon Michaud, Samuel Faucher, François Grondin, Jean-Samuel Lauzon, Mathieu Labbé, Dominic Létourneau, François Ferland, and François Michaud. 3d localization of a sound source using mobile microphone arrays referenced by slam. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10402–10407. IEEE, 2020.
- [62] Alvin Miranda, Joshua Vander Hook, and Christoph Schaal. Lamb wave-based mapping of plate structures via frontier exploration. *Ultrasonics*, 110:106282, 2021.
- [63] Amir Mobarhani, Shaghayegh Nazari, Amir H. Tamjidi, and Hamid D. Taghirad. Histogram based frontier exploration. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1128–1133, 2011.
- [64] Michael Montemerlo, Sebastian Thrun, Daphne Koller, Ben Wegbreit, et al. Fastslam: A factored solution to the simultaneous localization and mapping problem. *Aaai/iaai*, 593598, 2002.
- [65] Hans Moravec and Alberto Elfes. High resolution maps from wide angle sonar. In *Proceedings. 1985 IEEE international conference on robotics and automation*, volume 2, pages 116–121. IEEE, 1985.
- [66] José Neira and Juan D Tardós. Data association in stochastic mapping using the joint compatibility test. *IEEE Transactions on robotics and automation*, 17(6):890–897, 2001.
- [67] Nicholas O'Donoghue, Joel B Harley, Chang Liu, José MF Moura, and Irving Oppenheim. Maximum likelihood defect localization in a pipe using guided acoustic waves. In *2012 Conference Record of the Forty Sixth Asilomar Conference on Signals, Systems and Computers (ASILOMAR)*, pages 1863–1867. IEEE, 2012.
- [68] Amit Kumar Pandey, K Madhava Krishna, and Mainak Nath. Feature based occupancy grid maps for sonar based safe-mapping. In *IJCAI*, page 2172, 2007.
- [69] Fangrong Peng, Tiexing Wang, and Biao Chen. Room shape reconstruction with a single mobile acoustic sensor. In *2015 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 1116–1120. IEEE, 2015.
- [70] Cedric Pradalier, Othmane-Latif Ouabi, Pascal Pomarede, and Jan Steckel. On-plate localization and mapping for an inspection robot using ultrasonic guided waves: a proof of concept. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5045–5050. IEEE, 2020.
- [71] Swadhin Pradhan, Ghufraan Baig, Wenguang Mao, Lili Qiu, Guohai Chen, and Bo Yang. Smartphone-based acoustic indoor space mapping. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(2):1–26, 2018.
- [72] Kaizhi Qian, Yang Zhang, Shiyu Chang, Xuesong Yang, Dinei Florencio, and Mark Hasegawa-Johnson. Deep learning based speech beamforming. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5389–5393. IEEE, 2018.

- [73] Jianxi Qiu, Fucal Li, Saqlain Abbas, and Yanping Zhu. A baseline-free damage detection approach based on distance compensation of guided waves. *Journal of Low Frequency Noise, Vibration and Active Control*, 38(3-4):1132–1148, 2019.
- [74] N Quaegebeur, P Masson, D Langlois-Demers, and P Micheau. Dispersion-based imaging for structural health monitoring using sparse and compact arrays. *Smart Materials and Structures*, 20(2):025005, 2011.
- [75] Md Atiqur Rahman and Yang Wang. Optimizing intersection-over-union in deep neural networks for image segmentation. In *International symposium on visual computing*, pages 234–244. Springer, 2016.
- [76] David Ribas, Pere Ridao, Juan Domingo Tardós, and José Neira. Underwater slam in a marina environment. In *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1455–1460. IEEE, 2007.
- [77] Ayoub Ridani, Othmane-Latif Ouabi, Nico F Declercq, and Cédric Pradalier. On-plate autonomous exploration for an inspection robot using ultrasonic guided waves. In *2021 European Conference on Mobile Robots (ECMR)*. IEEE, 2021.
- [78] Daniele Salvati, Carlo Drioli, and Gian Luca Foresti. Sound source and microphone localization from acoustic impulse responses. *IEEE Signal Processing Letters*, 23(10):1459–1463, 2016.
- [79] Usama Saqib and Jesper Rindom Jensen. A model-based approach to acoustic reflector localization with a robotic platform. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4499–4504. IEEE, 2020.
- [80] Usama Saqib and Jesper Rindom Jensen. A framework for spatial map generation using acoustic echoes for robotic platforms. *Robotics and Autonomous Systems*, page 104009, 2022.
- [81] AG Sazontov and AI Malekhanov. Matched field signal processing in underwater sound channels. *Acoustical Physics*, 61(2):213–230, 2015.
- [82] Christoph Schaal, Matthew Brown, and Katja Schulz. Experimental investigation of lamb wave-based edge detection methods. In *Health Monitoring of Structural and Biological Systems XIII*, volume 10972, pages 440–450. SPIE, 2019.
- [83] Novonil Sen and Tribikram Kundu. Acoustic source localization in a highly anisotropic plate with unknown orientation of its axes of symmetry and material properties with numerical verification. *Ultrasonics*, 100:105977, 2020.
- [84] Xiyu Song, Mei Wang, Hongbing Qiu, and Liyan Luo. Indoor pedestrian self-positioning based on image acoustic source impulse using a sensor-rich smartphone. *Sensors*, 18(12):4143, 2018.
- [85] Hans Steyskal and John F Rose. Digital beamforming for radar systems. *Microwave Journal*, 32:121–123, 1989.
- [86] Zhongqing Su and Lin Ye. *Identification of damage using Lamb waves: from fundamentals to applications*, volume 48. Springer Science & Business Media, 2009.

-
- [87] Zhongqing Su, Lin Ye, and Ye Lu. Guided lamb waves for identification of damage in composite structures: A review. *Journal of Sound and Vibration*, 295(3):753 – 780, 2006.
- [88] Hu Sun, Aijia Zhang, Yishou Wang, and Xinlin Qing. Baseline-free damage imaging for metal and composite plate-type structures based on similar paths. *International Journal of Distributed Sensor Networks*, 15(4):1550147719843054, 2019.
- [89] Johan Fredrik Synnevag, Andreas Austeng, and Sverre Holm. Adaptive beamforming applied to medical ultrasound imaging. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 54(8):1606–1613, 2007.
- [90] Morteza Tabatabaeipour, Oksana Trushkevych, Gordon Dobie, Rachel S Edwards, Steven Dixon, Charles MacLeod, Anthony Gachagan, and S Gareth Pierce. A feasibility study on guided wave-based robotic mapping. In *2019 IEEE International Ultrasonics Symposium (IUS)*, pages 1567–1570. IEEE, 2019.
- [91] Morteza Tabatabaeipour, Oksana Trushkevych, Gordon Dobie, Rachel S Edwards, Ross McMillan, Charles Macleod, Richard O’Leary, Steve Dixon, Anthony Gachagan, and Stephen G Pierce. Application of ultrasonic guided waves to robotic occupancy grid mapping. *Mechanical Systems and Signal Processing*, 163:108151, 2022.
- [92] S. Theodoridis, R. Chellappa, M. Viberg, and A. Zoubir. *Academic Press Library in Signal Processing: Array and Statistical Signal Processing*. ISSN. Elsevier Science, 2013.
- [93] Sebastian Thrun. Learning occupancy grid maps with forward sensor models. *Autonomous robots*, 15(2):111–127, 2003.
- [94] Sebastian Thrun, Wolfram Burgard, Dieter Fox, et al. *Probabilistic robotics*. MIT press Cambridge, 2005.
- [95] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- [96] Guanyu Tube. <https://tubingchina.com/Ultrasonic-Velocity-Acoustical-Properties-Of-Common-Materials.htm>. Accessed: 2022-06-21.
- [97] Barry D Van Veen and Kevin M Buckley. Beamforming: A versatile approach to spatial filtering. *IEEE assp magazine*, 5(2):4–24, 1988.
- [98] Arnoud Visser and Bayu A Slamet. Balancing the information gain against the movement cost for multi-robot frontier exploration. In *European Robotics Symposium 2008*, pages 43–52. Springer, 2008.
- [99] Wenyi Wang, Thomas C Henderson, Anshul Joshi, and Edward Grant. Slambot: Structural health monitoring robot using lamb waves. In *2014 International Conference on Multisensor Fusion and Information Integration for Intelligent Systems (MFI)*, pages 1–7. IEEE, 2014.
- [100] Rob Worley, Yicheng Yu, and Sean Anderson. Acoustic echo-localization for pipe inspection robots. In *2020 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, pages 160–165. IEEE, 2020.

- [101] Kai M Wurm, Cyrill Stachniss, and Giorgio Grisetti. Bridging the gap between feature-and grid-based slam. *Robotics and Autonomous Systems*, 58(2):140–148, 2010.
- [102] Brian Yamauchi. A frontier-based approach for autonomous exploration. In *Proceedings 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation CIRA '97. 'Towards New Computational Principles for Robotics and Automation'*, pages 146–151. IEEE, 1997.
- [103] Shenxin Yin, Zhiwen Cui, and Tribikram Kundu. Acoustic source localization in anisotropic plates with “z” shaped sensor clusters. *Ultrasonics*, 84:34–37, 2018.
- [104] Paul Zabbal, Guillemette Ribay, Bastien Chapuis, and Julien Jumel. Multichannel multiple signal classification for dispersion curves extraction of ultrasonic guided waves. *The Journal of the Acoustical Society of America*, 143(2):EL87–EL92, 2018.
- [105] Liang Zeng, Liping Huang, Zhi Luo, and Jing Lin. Damage imaging that exploits multipath scattered lamb waves. *Structural Health Monitoring*, 19(6):1629–1644, 2020.
- [106] Hanyu Zhan, Hanwan Jiang, and Ruinian Jiang. Three-dimensional images generated from diffuse ultrasound wave: detections of multiple cracks in concrete structures. *Structural Health Monitoring*, 19(1):12–25, 2020.
- [107] Jucheng Zhang, Yunfeng Han, Cuie Zheng, and Dajun Sun. Underwater target localization using long baseline positioning system. *Applied Acoustics*, 111:129–134, 2016.
- [108] Meijie Zhao, Wensong Zhou, Yong Huang, and Hui Li. Sparse bayesian learning approach for propagation distance recognition and damage localization in plate-like structures using guided waves. *Structural Health Monitoring*, 20(1):3–24, 2021.