



Subject Areas:

...

Keywords:

Ultrasound tomography; Guided waves; Inversion; Thickness mapping

Author for correspondence:

P. Huthwaite

e-mail: p.huthwaite@imperial.ac.uk

Evaluation of inversion approaches for guided wave thickness mapping

P. Huthwaite

Department of Mechanical Engineering, Imperial College, London, SW7 2AZ, United Kingdom

Accurate inversion is vital for quantitative imaging, particularly ultrasonic guided wave tomography, where thickness maps of plate-like structures are reconstructed to quantify corrosion damage. The dispersive properties of guided waves are often exploited to enable thickness maps to be produced from wave speed reconstructions. Ray tomography, diffraction tomography, and a hybrid algorithm combining their features were investigated to reconstruct wave speed. Test data produced from simple defects of different sizes using (1) a realistic full elastic guided wave model, and (2) the equivalent idealised acoustic model, was passed to the imaging algorithms, generating wave speed maps, and from these, thickness maps. For both data sets, ray tomography exhibited poor resolution. Diffraction tomography performed better, but was limited to shallow, small defects. The hybrid algorithm achieved the best results, giving a resolution around 1.5-2 wavelengths from test data (1) compared to half wavelength from (2). These results were validated with experimental data, and also extended to a realistic corrosion patch confirming the trends demonstrated with simple defects. The resolution loss with realistic test data (1) compared to idealised data (2) indicates the acoustic model cannot accurately capture guided wave scattering and an alternative approach is necessary for better resolution reconstructions.

1. Introduction

Quantitative imaging, i.e. the reconstruction of maps of properties through a domain from sets of measurements, has applications in a number of areas including medicine [1,2] and geophysics [3]. One important application of quantitative imaging methods is guided wave tomography. For this, Lamb-type guided waves are excited in a plate, transmitted across a domain, and then measured at several points by an array; these measurements are then used to produce a map of plate thickness across the domain. Determining the remaining wall thickness of vessels and pipes is very important for corrosion detection and quantification in a number of industries, particularly petrochemical. Conventional ultrasonic thickness maps are produced by scanning a probe across all points on the surface, and obtaining the thickness from the arrival time of the wave reflected from the backwall. In contrast, guided wave tomography removes the requirement to have direct access to all points on the surface, and is also faster since the entire surface need not be scanned. In many cases any curvature is sufficiently low that they can be approximated as flat plates, so the standard Lamb wave dispersion relations can be used.

The process of quantitative imaging can be considered an inversion problem, to determine what particular map of properties would have caused a given set of measurements. A significant challenge of quantitative imaging is to identify a suitable forward model to describe the behaviour of the wave propagation. An ideal scenario would be a model which is (a) accurate at describing the physics of interest, (b) highly sensitive to changes in the properties to be reconstructed, (c) insensitive to changes in any unwanted/irrelevant properties, (d) relatively fast, and (e) easily invertible. The last requirement can take a number of forms; some forward models can be directly inverted (e.g. [4,5]), while some may require an iterative, gradient stepping approach (e.g. [6, 7]), which would place a strong emphasis on models which contain few local minima and have straightforward methods of gradient calculation. The ultimate speed of the algorithm will be a combination of the speed of the forward model and the complexity of its inversion. While the focus of this paper is on inverting mechanical wavefields, similar research has been ongoing for the inversion of electrical impedance measurements [8] and electromagnetic waves [9].

The dispersive nature of guided waves in plates provides a widely exploited approach to simplify the inversion. Lamb waves at a certain frequency will have particular known phase and group velocities for an infinite plate of particular thickness. Based on this, a commonly used assumption in guided wave tomography inversion is that a guided wave travelling in a plate of varying thickness and dimensions x and y in-plane and z out-of-plane will behave in the same way as an acoustic wave travelling in a medium across the x and y dimensions with varying velocity, where the velocity at each location is given by the velocity of the Lamb wave at that frequency and thickness. This is a very valuable assumption for greatly simplifying the inversion process, enabling the full 3D elastic problem to be recast as a 2D acoustic problem, and has been widely demonstrated in the production of thickness maps [10,11]. However, the effects of this assumption on the quality of the inversion is unclear; no attempts have been made to quantify the errors or to understand the circumstances where the assumption fails.

The ray theory of geometric optics is another widely used assumption, and can be used to invert an acoustic wavefield. The waves are assumed to propagate as rays, with the arrival time being a line integral of the slowness (reciprocal of the velocity) along the ray path. Variations exist: the rays can be assumed straight [12,13], or can be bent to account for refraction as they pass through the domain [14,15]. However, the act of ignoring diffraction causes the reconstruction to be severely resolution limited for the majority of applications [16]. Despite this, the model has widespread use in guided wave tomography reconstructions [10,17,18].

Higher resolution attempts have been made using diffraction tomography [19], which uses the Born approximation as the forward assumption. This, however, relies on the object being weakly scattering, so any defect imaged must be small and of low contrast (i.e. very shallow), which is unlikely to be the case for the majority of defects of interest. The Rytov approximation has been considered for this problem [20] but the main challenge is that it requires the phase of the field

to be unwrapped, which is difficult for complex scatterers in the presence of noise. HARBUT (the hybrid algorithm for robust breast ultrasound tomography) [21] uses a low resolution ray tomography background to reduce the effective contrast of any remaining scatterers, which can then be imaged using a Born approximation approach, and has been successfully applied to invert breast ultrasound tomography data [22]. The approach was applied to the acoustic inversion stage of guided wave tomography in [11], enabling the higher resolution associated with diffraction tomography to be achieved across a wider variety of defects. While for purely acoustic problems, it is known that the resolution is limited to the diffraction limit $\lambda/2$, it is unclear what the resolution performance is when test data is produced from the realistic fully 3D guided wave model.

This paper aims to comprehensively and quantitatively evaluate the accuracy of the underlying assumptions used in guided wave tomography, and to understand the physical mechanisms causing these to fail. Sec. 2 outlines the background theory for the guided wave propagation and the imaging algorithms. Initial investigations outlined in Sec. 3 reconstruct sets of simple defects across a range of sizes and depths, with the results being validated with some experimental tests. Section 4 investigates imaging with more complex defects to confirm whether the results generalise, and Sec. 5 discusses the results and their implications.

2. Background

(a) Guided wave theory

Lamb waves [23] are guided elastic waves which travel in infinite plates of constant thickness. As with all guided waves, they can be considered as a superposition of bulk waves (longitudinal and shear) combined using particular boundary conditions to form a particular characteristic mode shape. For Lamb waves, the mode shapes of displacement and stress can be calculated across the thickness of the plate, to fit the ‘sound soft’ (i.e. zero pressure) free boundaries at the top and bottom surfaces. Several solutions can be found to these equations depending on the frequency and thickness, and software packages such as Disperse [24,25] exist to calculate these. Figure 1(a) plots the mode shapes for the fundamental A_0 and S_0 modes, which exist at low frequencies. Figures 1(b) and (c) plot the phase velocity and group velocity dispersion curves respectively as a function of the product of plate thickness and frequency.

The Lamb wave formulation assumes that the plate is of uniform thickness, which in general is not the case for guided wave tomography since the focus is on estimating varying thickness. However, provided that the variation in thickness is slowly varying, it can be assumed that one mode behaves as a Lamb wave with the acoustic properties (phase and group velocities) associated with the local thickness. Under this assumption, a 2D acoustic bulk wave model can be defined with equivalent velocities which would approximate the behaviour of the guided waves. An aim of this paper is to identify how rapidly the thickness can vary before the assumption becomes invalid.

Since the guided waves are dispersive, the acoustic wave equation is expressed in the frequency domain as

$$\rho \left[\frac{\partial}{\partial x} \left(\frac{1}{\rho} \frac{\partial p}{\partial x} \right) + \frac{\partial}{\partial y} \left(\frac{1}{\rho} \frac{\partial p}{\partial y} \right) \right] + k^2 p = 0 \quad (2.1)$$

where p represents the Fourier transform of the pressure (which could be substituted for some parameter of the guided wave mode, such as the out-of-plane displacement), $k = \omega/c$ is the local wavenumber in terms of the angular frequency ω and the phase velocity c , and ρ is the density. For the purposes of this paper, ρ is assumed constant since its effect on the wavefield tends to be negligible compared to the that of the velocity variations (although it could be set proportional to the plate thickness); it therefore cancels out of the equation, which can be rewritten as

$$(\nabla^2 + k^2) p = 0 \quad (2.2)$$




Figure 1. Lamb waves in a steel plate. (a) presents the mode shapes for the A_0 and S_0 modes, at 0.5 MHz.mm and 2 MHz.mm respectively. The plots show the displacement of points in a length of the plate at a single point in time. (b) gives the phase velocity dispersion curves, and (c) gives the group velocity.

which is the Helmholtz equation. Having made the assumption that the waves follow this model, this equation itself must be inverted via an acoustic inversion technique.

(b) Imaging algorithms

Acoustic velocity reconstruction algorithms are the primary algorithms used in guided wave tomography, and are thus dependent on the accuracy of the mapping between elastic guided waves and acoustic waves discussed above. This paper will consider three of these algorithms, which each comprise two core features: a set of physical assumptions which are used to describe how the waves propagate through the domain and a mathematical inversion method, which is used to calculate the best solution which fits the data. These three algorithms cover the majority of the acoustic wave speed inversion approaches used, and are the focus of this paper since they have been demonstrated to be fast, robust and practical for a wide range of configurations, but it is recognised that alternative approaches exist such as full wave inversion [26] and the Novikov algorithm [27].

(i) Ray tomography

The first algorithm uses the ray assumption, which neglects diffraction. Diffraction is a phenomenon which occurs when a wavefield interacts with an object that is small relative to the wavelength. If a situation exists where the effects of diffraction can be neglected, such as for large scatterers, then the ray theory of geometrical optics will be suitable to describe the wave propagation. By considering a solution of the form $p = Ae^{-i\omega\tau(x,y)}$, where τ represents a time delay and A is an arbitrary complex constant, it is possible to rewrite eq. (2.2), as

$$\begin{aligned} k^2 Ae^{-i\omega\tau} &= -\nabla^2 (Ae^{-i\omega\tau}) \\ &= -\nabla \left(\nabla Ae^{-i\omega\tau} - i\omega \nabla \tau Ae^{-i\omega\tau} \right) \\ &= -\nabla^2 Ae^{-i\omega\tau} + 2i\omega \nabla \tau \nabla Ae^{-i\omega\tau} + i\omega \nabla^2 \tau Ae^{-i\omega\tau} + A\omega^2 (\nabla \tau)^2 e^{-i\omega\tau}. \end{aligned} \quad (2.3)$$

By recognising that increased frequency causes shorter wavelengths, making any object become large relative to the wavelength, it is possible to remove any components associated with diffraction by considering the asymptotic limit of this equation as ω tends to infinity,

$$k^2 = \omega^2 (\nabla \tau)^2 \quad (2.4)$$

where any terms not containing orders of less than ω^2 (note that $k^2 = \omega^2/c^2$) have been dropped, and both sides have been divided by $Ae^{-i\omega\tau}$ to remove the common factor. The equation can be rearranged to become

$$\frac{1}{c} = |\nabla \tau| \quad (2.5)$$

which is the eikonal equation [28], and is a valid representation of the wavefield assuming that any object or variation within that field is large relative to the wavelength. The product $\omega\tau$ in the exponential term represents the (unwrapped) phase of the wavefield, so τ corresponds to the 'phase' arrival time. Since ray tomography reconstructions are typically based on the arrival times of the wavepackets τ_{gp} , the equation is altered to use group velocity c_{gp}

$$\frac{1}{c_{gp}} = |\nabla \tau_{gp}|. \quad (2.6)$$

An eikonal solver, such as the fast marching method [29,30], uses a finite difference approximation to calculate $\tau_{gp}(x, y)$ on a uniform grid representing the domain. The inversion process of ray

tomography then involves determining a suitable group velocity field to fit this model to a set of arrival times extracted from measurements. The initial step is therefore to extract the arrival times from the experimental time traces, and a wide variety of arrival time pickers have been developed to achieve this, with developments mainly made in the field of geophysics. These use ideas such as cross correlation [31,32], the Akaike information criterion [33], and wavelet transforms [34] but the challenges of determining accurate arrival times mean that manual arrival time picking is still common (see for example [26]). In this paper, in the reconstructions from simulated data, which is very clean, a thresholding approach is used. The arrival time is taken when the Hilbert envelope of the signal exceeds 5% of the signal peak. For the experimental data sets, where noise was present, the arrival times were manually picked.

Having obtained the arrival times, a technique must then be developed which produces a map of wave velocity to match these. In many inversion algorithms the approach is to locally ‘linearise’ the model of wave propagation in some way. In some cases the solution can then be calculated by a direct inversion, if the linear model accurately describes the wave behaviour between the starting point (typically a homogeneous field) and the solution. Alternatively, if the linearity only holds locally, many small steps can be taken, with the gradient recalculated at each step, which is an iterative gradient-type approach. It is noted that alternative approaches to determining a global minimum such as evolutionary methods [35,36] have been applied to imaging inversions but these have not found widespread use.

Both the direct and gradient-based approaches can be applied to ray tomography. The eikonal equation (2.6) can be rewritten as an integral of ‘slowness’ (reciprocal of speed) along the ray path $(x(r), y(r))$

$$\tau_{gp} = \int \frac{1}{c_{gp}(x(r), y(r))} dr. \quad (2.7)$$

If the ray path is fixed, for example if straight ray paths are assumed, then τ_{gp} can be seen to be linear with respect to slowness, due to the linearity of the integral operator. This enables it to be inverted using a direct approach, such as the filtered backprojection method, widely used in practice for reconstructing attenuation maps from the amplitudes of x-ray projections [37], which is an equivalent problem.

For more accurate reconstructions, refraction should be included by bending the ray paths; however this makes the problem non-linear, so it can no longer be solved with a single direct step. This paper uses this more accurate model, with the conjugate gradient method exploited to perform the inversion. This is a common approach and full details of the implementation are available elsewhere [38,39]. In summary, a cost function is defined as

$$C = \sum_{i=1}^{N \times N} \left(\tau_{model}^i - \tau_{meas}^i \right)^2 \quad (2.8)$$

where the sum is performed across all the $N \times N$ different send-receive pairs; the goal is to fit τ_{model} (the arrival times modelled with the eikonal solver) to τ_{meas} (the measured arrival times) by minimising C , which is achieved by a gradient-based stepping approach. For this, the gradient of C for variations in the modelled slowness field $s = 1/c_{gp}$ must be calculated. A discretised form of s is considered, which consists of M tiles with independent slowness values, given by s_m , for $m = 1..M$, to represent the slowness field. For the change in C with the slowness in tile m

$$\frac{\partial C}{\partial s_m} = \sum_{i=1}^{N \times N} \frac{\partial \left(\tau_{model}^i - \tau_{meas}^i \right)^2}{\partial s_m} \quad (2.9)$$

$$= \sum_{i=1}^{N \times N} 2 \left(\tau_{model}^i - \tau_{meas}^i \right) \frac{\partial \tau_{model}^i}{\partial s_m}. \quad (2.10)$$

Taking the length of the ray i within the tile m as $l_{i,m}$, by inspection of eq. (2.7) the gradient can be written as

$$\frac{\partial \tau_{model}^i}{\partial s_m} = l_{i,m} \quad (2.11)$$

so

$$\frac{\partial C}{\partial s_m} = \sum_{i=1}^{N \times N} 2 \left(\tau_{model}^i - \tau_{meas}^i \right) l_{i,m}. \quad (2.12)$$

This makes the gradient very straightforward to calculate; the process is to calculate the residual $\tau_{model}^i - \tau_{meas}^i$ for each ray path, then double this and project it along the ray path, adding it to all the tiles it intersects according to the length of the ray path within each tile. This gradient is not used directly, but forms the basis of the conjugate gradient stepping method, with full details described in [39].

The resolution limit, defined in this paper as the wavelength of the maximum wavenumber component accurately reconstructed in the image, is dependent on exactly how the arrival times are picked. If a cross-correlation approach is used for this then the resolution is close to the width of the first Fresnel zone $\sqrt{L\lambda}$, where L is the distance between the furthest two transducers in the array and λ is the wavelength [16,40]. Reference [41] provides a physical explanation of the causes of this resolution limit; in summary, diffraction causes ‘wavefront healing’ to occur, which smooths the wavefronts once they have interacted with any features and hence reduces the resolution of any reconstruction from the arrival times. By focusing on arrival time pickers which extract the earliest arrivals rather than the centre of the wavepacket (as is the case with the cross-correlation approach), the result will be less affected by the diffraction-type scattering and it may be possible to improve on this resolution limit to an extent. Using the ray model to describe wave propagation is very convenient for inversion, despite the limited resolution, because there are very few local minima. This is straightforward to understand conceptually since one would expect the term $\tau_{meas} - \tau_{model}$ to be monotonic with changes in the contrast and size of the scatterer.

It should be noted that the Rytov approximation [37] employs a similar approach to the eikonal model by expressing the field as a phase; this approach is not considered here since the algorithm requires phase unwrapping which can be challenging given practical issues such as noise and limited spatial sampling.

(ii) Diffraction Tomography

Diffraction tomography commonly utilises the first Born approximation to produce a reconstruction [37,42], although the Rytov approximation discussed above can also be used. Since the Born approximation accounts for diffraction, the resolution limit is significantly improved over the ray tomography approaches. The Helmholtz equation, eq. (2.2) can be reformulated as

$$p(x, y) = p_0(x, y) + \iint p(x', y') G(x', y', x, y) O(x', y') dx' dy' \quad (2.13)$$

where p_0 is the incident field (i.e. what would be measured if no scatterer were present), G is the free space Green’s function, and O is the object function, defined as

$$O(x, y) = k_0 \left[\left(\frac{c_0}{c(x, y)} \right)^2 - 1 \right] \quad (2.14)$$

with k_0 and c_0 being the background wavenumber and phase velocity respectively.

The first Born approximation makes the assumption that the scatterer is weakly scattering, i.e. it is small, and of low contrast. A criterion for this assumption being acceptably accurate is that the phase shift through the scatterer (which causes the difference between p and p_0) must be small, less than a quarter of a cycle [37]. Under this, as a first order assumption, the difference

between p and p_0 will be small, so the p term within the integral can be replaced with p_0 leading to

$$p = p_0 + \iint p_0 G O dx' dy'. \quad (2.15)$$

As with ray tomography, this has linearised the equation, so that $p - p_0$ varies linearly with O , and as before a direct solution can be obtained. To analyse this more, if the sources and receivers are in the far field such that the paraxial approximation holds [42,43] (although it should be noted that this is not a necessary condition for diffraction tomography algorithms), both the incident field and the Green's functions can be expressed as plane waves, so

$$p_s(\hat{\mathbf{s}}_0, \hat{\mathbf{s}}) = \int \exp(ik_0 \hat{\mathbf{s}}_0 \cdot \mathbf{x}') \exp(-ik_0 \hat{\mathbf{s}} \cdot \mathbf{x}') O(\mathbf{x}') d\mathbf{x}' \quad (2.16)$$

$$= \int \exp[ik_0(\hat{\mathbf{s}}_0 - \hat{\mathbf{s}}) \cdot \mathbf{x}'] O(\mathbf{x}') d\mathbf{x}'. \quad (2.17)$$

where $p_s(\hat{\mathbf{s}}_0, \hat{\mathbf{s}})$ is the scattered field (i.e. $p - p_s$) for source and receiver directions given by the unit vectors $\hat{\mathbf{s}}_0$ and $\hat{\mathbf{s}}$ respectively, and the double integral across the single dimensions x' and y' has been replaced with a single integral over the 2D space $\mathbf{x}'$. By expressing the 2D Fourier transform of O as $\hat{O}$, then

$$p_s(\hat{\mathbf{s}}_0, \hat{\mathbf{s}}) = \hat{O}[-k_0(\hat{\mathbf{s}}_0 - \hat{\mathbf{s}})] \quad (2.18)$$

i.e. under the Born approximation in the far field (although the results generalise outside the paraxial approximation), each component of the scattered field can be mapped to a particular component of the Fourier transform of the object function. This has important implications for the resolution of the algorithm, and will be exploited later. The highest spatial frequency component of $\hat{O}$ measurable from p_s appears when $\hat{\mathbf{s}} = -\hat{\mathbf{s}}_0$, and is $2k_0$; this means that the wavefield can only encode a low-pass filtered version of O , making the well known diffraction limit for resolution $\lambda/2$.

(iii) HARBUT

The limitation of the Born approximation to small, weakly scattering objects contrasts to the ray tomography approach, which can only reconstruct the large features where diffraction is limited. HARBUT aims to combine the complementary features of these two algorithms, to enable high resolution reconstructions for a range of scattering sizes and contrasts, and was introduced for breast ultrasound tomography in [21] and applied to guided wave tomography in [11]. The HARBUT algorithm decomposes the object function into two components

$$O = O_b + O_\delta. \quad (2.19)$$

In this, the background object function O_b is a smoothly varying approximation of O and O_δ is the remainder. O_b can be obtained from the ray tomography reconstruction, and eq. (2.15) can be rewritten as

$$p = p_b + \iint p_b G_b O_\delta dx' dy' \quad (2.20)$$

where p_b and G_b are estimates of the incident wavefield and Green's functions respectively in the background O_b ; full details of how they are calculated are in [21]. As before, this forms a linearisation, provided p_b and G_b remain constant, which enables O_δ to be calculated directly.

Reference [11] recognised that, in many cases in guided wave tomography, reconstructions could be improved by doing additional iterations. This is achieved by setting $O_b^{(n)} = O^{(n-1)}$, with n being the iteration number, updating p_b and G_b accordingly, calculating the new $O_\delta^{(n)}$ and from it $O^{(n)}$. Instead of being a direct approach, or a gradient stepping approach, this iterative formulation follows the Newton-Raphson method. At each point, the problem is linearised




Figure 2. (a) Configuration of the simple, axisymmetric defect setup. Due to symmetry a single source is excited and measurements are taken at the locations marked, then the different send-receive pairs for the full array can be generated by copying this data. A cross section through the defect is shown in (b), with the D and W parameters marked.

against the background O_b , and O_δ is the solution to the problem under that linearisation. The linearisation is then updated for this new point and the process repeats. By iterating, we are accounting for any non-linearities there are in the model around the solution point, which allows a greater range of sizes and defect depths to be accurately reconstructed. The value of using the ray tomography image as a starting point for iterative HARBUT, rather than just using a homogeneous background is to minimise the likelihood of convergence to a local minimum. The theoretical resolution of HARBUT and iterative HARBUT maintain that of diffraction tomography, $\lambda/2$.

3. Simple defect analysis

This section aims to present a broad picture of the performance and implications of the different assumptions, by reconstructing the thicknesses of a series of different sized defects. A simple axisymmetric defect shape is chosen, allowing each defect to be characterised purely by its outer diameter and its peak depth. To produce the necessary test data for the different defects in this study, simulations must be performed. Given the potential for confusion between these simulations and the forward models which describe the wave propagation within the algorithms themselves, a naming convention will be used throughout this paper. The simulated experimental test data will be referred to as ‘type 1’ and ‘type 2’ data depending on how they are produced, as outlined below. By contrast, ray tomography, DT and HARBUT which are used to subsequently produce the reconstruction from the data will be referred to as ‘acoustic imaging algorithms’.

(a) Configuration

A diagram of the configuration is presented in Fig. 2(a). The plate is made from 10mm thick steel (although it should be noted that the results are presented normalised, so are valid for any plate thickness provided that the waves are excited at the same frequency-thickness point), and an axisymmetric defect is placed at the centre of the plate. A circular array of radius 180mm is positioned coaxially, with $N = 64$ transducers; in theory $N \times N$ time traces should be recorded through the different send-receive pairs to enable all the information to be collected. However, due to symmetry, just a single source may be used for illumination; the data from this can then be replicated to produce a full matrix of data. Symmetry can also be exploited to generate a set of measurements across all 360 degrees from the source by taking measurements at only half the angles (marked in grey in Fig. 2(a)) and mirroring the results. A five-cycle Hann-windowed toneburst, at a centre frequency of 50kHz, was used to excite an A_0 wave at the source. This gives a wavelength of around 37mm at the centre frequency for the 10mm plate.




Figure 3. (a) The meshing approach used in the finite element models (N.B. not to scale), used to produce experimental test data from the elastic model, type 1 data. (b) presents the finite difference approach for producing type 2 experimental data, using the acoustic wave equation. Inset is the a zoomed section of the finite difference grid, showing the nodes; each of these nodes has a velocity associated with it corresponding to the underlying thickness map, as shown by different shades of grey. This example shows the velocity at 50 kHz.

The defect cross-section needs to be selected; this can then be rotated to sweep out the axisymmetric shape. The initial focus is on simple defects, so a raised cosine Hann function is used

$$T(r) = \begin{cases} T_0 - \frac{D}{2} \left[1 + \cos\left(\frac{2\pi r}{W}\right) \right] & r < W/2 \\ T_0 & r \geq W/2 \end{cases} \quad (3.1)$$

with W defining the outer diameter of the defect, D defining the depth, T_0 representing the nominal plate thickness and r being the radius from the centre. This function has a smoothly varying shape, and the two parameters D and W , illustrated in Fig. 2(b) can be used to vary the dimensions of the defect to enable a variety of cases to be tested.

(b) Elastic Model

Various solution approaches exist for simulating ultrasonic waves in elastic media. For simulating the configuration outlined above, the finite element (FE) method was selected since it is flexible enough to simulate a variety of complex shapes; throughout this paper, following the convention above, the test data produced from this model will be referred to as ‘type 1’ data. This is solved in the time domain using explicit finite difference steps; this explicit formulation allows larger models to be simulated since no memory demanding matrix inversion steps need to take place. The majority of the models were solved with the open source Pogo software package [44], a finite element solver which runs very quickly on Nvidia graphics cards. Since this software is relatively new, initial models were performed using both Pogo and the Abaqus Explicit package to validate the accuracy of the software, confirming that the results were the same.

The meshing approach used in the finite element models in this paper is illustrated in Fig. 3(a). A structured uniform grid of hexahedral brick elements is used to model the plate, then to model the thickness reduction caused by corrosion, the out-of-plane dimension of the elements is

reduced accordingly. This approach reduces the mesh scattering which occurs with free meshing, although there is likely to be a degree of anisotropy in the wave propagation due to the element compression, but this should be small provided the mesh is sufficiently refined. Typical estimates of the errors associated with this are available in [45]. Also, it is known that around the chosen operating point (0.5 MHz.mm with A_0) the mode shape does not change significantly as thickness reduces, so this meshing approach has the beneficial property that the same number of elements is available to capture the wavefield variation through the thickness.

For the models, first order linear hexahedral brick elements were used. In Abaqus these had reduced integration with hourglass control, while the Pogo elements were fully integrated. The elements were sized 1 mm along the two in-plane dimensions, and 10 elements were used through the plate thickness, making $1 \times 1 \times 1 \text{ mm}^3$ sized cubic elements in the 10mm thick region away from the defect. Given the wavelength of around 37mm for the A_0 mode at 50kHz, this gives 37 elements per wavelength, which is well above the typical 20-30 elements per wavelength considered necessary to accurately capture the behaviour of the elastic waves [45].

The plate was defined to be $560 \times 560 \text{ mm}^2$ in size. Absorbing boundaries were applied to the edge of the plate to avoid boundary reflections. The Absorbing Layers using Increasing Damping (ALID) technique outlined in [46] was used; the absorbing region extended 80mm from each boundary (leaving a 'working region' of $400 \times 400 \text{ mm}^2$) and the mass proportional damping factor was defined as

$$\alpha = Kl^3 \quad (3.2)$$

where l is the fraction of the distance through the layer, and K is a constant. From testing, $K = 2 \times 10^6$ was chosen to minimise the boundary reflections.

(c) Acoustic Model

An important aim of this paper is to distinguish where in the thickness reconstruction process the errors occur; of the simplifying assumptions, these could either be caused by the guided wave to acoustic mapping of Sec. (a), or by the acoustic imaging algorithms of Sec. (b) which are subsequently used to perform the reconstruction. To aid this distinction, an acoustic bulk wave model was run, with velocities chosen to match those of the plate based on the Lamb wave propagation; the test data produced will be referred to as 'type 2' data. The velocity map at 50 kHz for a defect 3.2λ wide and 60% deep is shown in Fig. 3(b). This model reproduces the same acoustic assumption as is used in the forward models of the imaging algorithms. The images generated from this data can be compared to the images generated from the type 1 data from the full elastic guided wave models. Since the velocities are frequency dependent, simulations are performed in the frequency domain, with a different velocity field at each frequency. At each frequency, the output from this will describe a transfer function between the source and receiver; this can be multiplied by an input source (typically the Fourier transform of an input signal) to determine the received signal. The inverse Fourier transform of all these frequencies can be taken to obtain the output in the time domain. To model the 50 kHz five-cycle Hann-windowed toneburst, 40 frequencies from 2.5 to 100 kHz were simulated; this captures the bandwidth of the toneburst without significant components being removed due to filtering.

As discussed previously, modelling the plate in 3D requires an approach to capture the complex shape; the solution was to use the finite element method. For the 2D acoustic model the features are instead represented by wave velocity variations, so there is no need to model complex boundaries and the solution can instead be calculated on a uniform grid; this grid is shown in the inset of Fig. 3(b). Therefore, the simpler finite difference method is used. This was written in Fortran 95, and consists of forming a sparse matrix equation by discretising the wave equation of eq. (2.1), and then using the Pardiso (Parallel Direct Solver) version available as part of the Intel MKL (Math Kernel Library) to calculate a solution to the matrix for a particular source.

In the model, perfectly matched layers (PMLs) of the form described in [47] were used to prevent artificial reflections appearing at the boundary of the domain, although since the




Figure 4. Cross sections through reconstructions of the simple defect, from type 2 data (i.e. from an acoustic approximation of the guided wave experiment). The left column of (a), (d), (g) and (j) is for 60% (of nominal thickness) deep defects, the middle column of (b), (e), (h) and (k) is for the 30% deep defects and the rightmost column of (c), (f), (i) and (l) contains the reconstructions for a 10% deep defects. The rows define the width of the defect. The first row (a)–(c) is for 6.4λ (240mm) wide defects, the second (d)–(f) is for 3.2λ (120mm) wide defects, the third (g)–(i) is for 1.6λ (60mm) wide defects and the final row (j)–(l) is for 0.8λ (30mm) wide defects.

simulation was undertaken at a single frequency the convolutional formulation described in the paper was unnecessary. The 1120×1120 nodes were arranged in a uniform grid, spaced by 0.5mm; this is large enough to model the complete array and the PMLs at the boundaries.

(d) Results

Initially, we consider the results from the reconstructions generated from type 2 data. The focus here is purely on the performance of the algorithms themselves, since both the simulated experimental data and the imaging algorithms use the acoustic assumption. Cross sections through the reconstructions from type 2 data are presented in Fig. 4 for defects of depths 10, 30 and 60% of the nominal plate thickness (1, 3 and 6mm), with widths of 15, 30, 60 and 120mm, which correspond to 0.8λ , 1.6λ , 3.2λ and 6.4λ .

Ray tomography is known to be resolution limited, and clearly performs poorly as the width of the defect is reduced, across all defect depths. The typical measure of resolution limit for ray tomography is given as $\sqrt{L\lambda}$ [16], which in this case is around 120mm, or 6.4λ ; it should be noted, however, that this is based on estimating arrival times with a cross-correlation approach, i.e. using the centre of the wavepacket. By picking the arrival times (as we do) based




Figure 5. Cross sections through reconstructions of the simple defect, using type 1 data (i.e. experimental data from a full elastic guided wave model). The left column of (a), (d), (g) and (j) is for 60% (of nominal plate thickness) deep defects, the middle column of (b), (e), (h) and (k) is for the 30% deep defects and the rightmost column of (c), (f), (i) and (l) contains the reconstructions for a 10% deep defects. The rows define the width of the defect. The first row (a)–(c) is for 6.4λ (240mm) wide defects, the second (d)–(f) is for 3.2λ (120mm) wide defects, the third (g)–(i) is for 1.6λ (60mm) wide defects and the final row (j)–(l) is for 0.8λ (30mm) wide defects.

on the first arrival of the signal, where the wavepacket will be less affected by diffraction-based scattering, it is possible that the resolution could be improved slightly. For the 6.4λ wide defects, ray tomography determines the thickness quite accurately, with a maximum deviation of around 0.5mm in Fig. 4(b). As the size of the defect is reduced to 3.2 wavelengths, the accuracy of the reconstruction suffers, reflecting the limited resolution of the algorithm. Further reductions in width result in reconstructions with little in common with the original thickness map, due to the domination of diffraction at such scales.

The diffraction tomography (DT) algorithm is limited by the validity of the Born approximation, which requires that the reconstructed objects must be sufficiently small and low contrast that the phase shift is low. Following this, there are several cases with the shallow, small defects, where the performance of DT matches the true thickness map well. However, Figs. 4(a), (b), (d), (g) show the DT reconstruction giving a significant deviation; this is a widely known result of the scatterer being too strong for the Born approximation to be valid [37]. The performance of the HARBUT algorithm is very good at almost all defect sizes. The largest deviations occur for the narrowest, deepest defects; for the worst case (Fig. 4(j)) the discrepancy is around 1.5mm (15% of the nominal thickness) at the deepest point; for all other cases the maximum error is less than 1mm (10%).

Having evaluated the performance of the acoustic algorithms by using test data obtained from the type 2 model, we now consider the images that are produced when the more realistic 3D elastic guided wave forward model is used to model the experiment, producing type 1 data. Figure 5 presents the cross sections for the same set of defects as before, with depths of 10, 30 and 60% of nominal plate thickness and widths of 0.8λ , 1.6λ , 3.2λ and 6.4λ .

The ray tomography reconstructions for the 6.4λ width defects noticeably overestimate the defect depth, which is not the case with the same reconstructions from Fig. 4. This suggests that the defect introduces an additional shift into the arrival time of the wavepacket in the type 1 data which is not captured by the acoustic model used in the inversion, although this has not affected the HARBUT or DT reconstructions, which are based on the phase and amplitude at a single frequency. Similar to the results in Fig. 4, the ray tomography algorithm is unable to generate accurate thickness maps for the defect 3.2λ wavelengths wide, although it should be noted that the overall shapes are more accurately reconstructed than before, which could give an indication that the signals from the elastic guided wave model contain a lower level of diffraction than from the equivalent acoustic model.

The DT reconstructions are limited by the maximum phase shift in the same way as when type 2 data was used, but when this shift is low, the Born approximation is an accurate representation of the wave propagation, so DT typically performs well. However, it is clear that, unlike the reconstructions from type 2 data, for the 1.6λ and 0.8λ widths, the defect depth is significantly underestimated. The HARBUT reconstructions, as was shown previously in Fig. 4, match the best throughout, but show a similar tendency to fail when the width is reduced to 1.6λ and below. This suggests that the resolution is limited when using acoustic reconstruction algorithms to produce images from type 1 data, preventing accurate reconstructions of small (in the in-plane dimensions) defects.

In guided wave tomography, the primary aim is to determine the minimum remnant wall thickness; therefore one metric for quantifying the performance of the algorithm is to take the error in the maximum defect depth, normalised against the plate thickness. However, this parameter makes no assessment of the quality of the shape reconstruction. For a single simple defect such as this, this is relatively unimportant, but when considering a more complex defect which might be formed from a number of superposed simple defects then the accuracy of the reconstructed shape of each is important to ensure that the correct peak is determined. Because of this, an alternative metric is defined, based on the root mean squared (RMS) average

$$e = \sqrt{\frac{\iint [T(x, y) - T^*(x, y)]^2 w(x, y) dx dy}{T_0^2 \iint w(x, y) dx dy}} \quad (3.3)$$

where $T(x, y)$ is the thickness of the original map at coordinates x and y , and T^* is the approximate thickness map produced by the reconstruction algorithm. w has been defined as a windowing function, defining the extent and weighting to apply to the error prior to integration. The selected choice is to define the window as

$$w(r) = \begin{cases} \frac{1}{2} [1 + \cos(\frac{\pi r}{A})] & r < A \\ 0 & r \geq A \end{cases} \quad (3.4)$$

with

$$r = \sqrt{x^2 + y^2} \quad (3.5)$$

where A defines the outer radius of the window. A is set to double the radius of the defect, but is limited to 180mm, the radius of the array, for the largest defects. The window has a peak at $r = 0$ and varies smoothly down to zero at $r = A$. This biases the error metric towards errors at the centre of the defect, where the deepest point of the defect is, while also measuring the shape of the defect away from this point. The window extends beyond the defect itself to measure how accurately the background is captured; this is important if two defects were superposed close together. To perform the calculation in practice, the integral is replaced by a discrete sum of the values across each pixel in the image.

This error metric will be zero in the case of a perfect reconstruction, and otherwise will represent the weighted RMS error as a fraction of the nominal plate thickness. The error metric is plotted in Fig. 6 for the reconstructions from type 2 data, and Fig. 7 for the reconstructions from more realistic type 1 data. In both figures, the trends, for HARBUT and ray tomography in particular, seem quite consistent between the three defect depths. The error for DT increases for deeper defects since the higher contrast causes the Born approximation to become violated quicker.

Plotting the results in this way enables clearer comparisons to be drawn between the reconstructions from type 1 data and type 2 data. Interestingly the local peak in the ray tomography reconstruction around 1.5λ is reduced slightly when reconstructing from type 1 data, compared to type 2 data, which suggests that the diffraction at the defect, which limits the ray assumption, is reduced. The acoustic model used to produce type 2 data cannot represent mode conversion at the defect, so all diffracted wavefield components appear in the predicted measurements; by contrast, the diffracted components in the type 1 data will be lower since some energy will be lost to the S_0 or SH_0 modes. Given that the reconstructions of such defects are known to be poor from ray tomography anyway, however (see Figs. 5(g)-(i)) it is unlikely that this improvement could be exploited.

There is a clear shoulder around defects of diameter 1λ which appears in the error curves for reconstructions from type 1 data (Fig. 7), which is not present for the type 2 data reconstruction (Fig. 6). This reduction in reconstruction accuracy for type 1 data (from the elastic guided wave simulation of the experiment) relative to type 2 is due to the inaccuracy of the elastic guided wave to acoustic mapping outlined earlier, since type 2 data is generated from an acoustic simulation of the experiment and hence does not rely on the mapping. This suggests that better reconstructions may be possible from type 1 data if a more accurate representation of the guided wave scattering was developed and incorporated into the inversion processes, rather than using the mapping.

Ultimately the main concern for users of the method is the maximum defect depth. As explained above, in order to capture this for more complex defects, the overall shape reconstruction must be reasonable, and once this is achieved it is important to study the accuracy of the maximum depth reconstruction. Therefore, in addition to the weighted average RMS error, e , defined above, a second parameter is defined as

$$e_d = \frac{|T_m - T_m^*|}{T_0} \quad (3.6)$$

where subscript m represents the minimum thickness of the original, T , and reconstructed, T^* , thickness maps. If the reconstruction obtains the maximum wall loss to within 10% of the plate thickness and the average reconstruction within 5%, then it is considered to be sufficiently accurate, i.e. if $e_d < 0.1$ and $e < 0.05$. This latter threshold is marked on Figs. 6 and 7. As an example, from the cross sections from type 2 data, in Fig. 4(h) (1.6λ wide and 30% deep) the ray tomography deepest value estimation is quite accurate, so $e_d < 0.1$, but the shape is a poor representation, so $e > 0.05$. By contrast DT and HARBUT both satisfy both criteria in this case. For the type 1 data, in Fig. 5(d) (3.2λ wide and 60% deep), again the ray tomography maximum depth assessment is within 10% of the total wall thickness, but again the shape is insufficiently well reconstructed to satisfy the shape requirement. DT clearly satisfies neither criterion in this example, but HARBUT satisfies both.

To apply these criteria, simulated experimental data was generated using type 1 and type 2 models, and reconstructions were generated at a matrix of all the different combinations of defect depths (0.5, 1, 2, 3, 4, 5, 6 mm, i.e. 5% to 60%) and defect widths (360, 240, 160, 120, 80, 60, 40, 30, 20, 15, 10, 5 mm, i.e. 9.6λ down to $\lambda/8$), and for each of these, e and e_d were calculated. These were interpolated to a 400×400 uniformly sampled grid, to enable the results to be plotted. The two criteria $e_d < 0.1$ and $e < 0.05$ were applied to the three different algorithms, ray tomography, DT and HARBUT, reconstructing from both type 1 and 2 data, and the results for when both are satisfied are shown in Figs. 8 and 9 respectively. It should be noted that the errors in reconstruction are likely to increase fairly monotonically with contrast; this is caused by the linear, or near linear,




Figure 6. Weighted RMS average error as a percentage of plate thickness, for reconstructions from type 2 data (i.e. data from an acoustic approximation of the guided wave experiment). Errors are given for defects of (a) 10%, (b) 30% and (c) 60% of the plate thickness. A threshold for an error of 10% of the plate thickness is marked on each graph.

nature of the algorithms. However, measures e and e_d both present error as a fraction of nominal plate thickness, rather than the depth of the defect, which causes the error for the shallower (i.e. lower contrast) defects to appear lower, and hence the plots show that all the algorithms typically perform well under the criteria presented for the shallower defects.

When reconstructing from both type 1 and 2 data, these plots demonstrate that the HARBUT algorithm performs well across all depths for most defect widths; for the shallower defects the resolution limit becomes apparent. The ray tomography reconstruction is resolution limited, so is only sufficiently accurate for larger defects. Diffraction tomography is limited in two ways, firstly by the fundamental resolution limit of the Born approximation, where it matches the HARBUT algorithm. As discussed earlier, the Born approximation also requires the phase shift through the object to be low; since the phase shift is a product of the size and the contrast (related to defect depth for guided wave tomography) of the object, an accurate reconstruction becomes a trade-off between these two parameters, producing the near-hyperbolic curve A marked in Fig. 8. To produce this curve, an approximation was made where the raised cosine Hann function from eq. (3.1) describing the defect shape was replaced with a simple flat-bottomed hole with a diameter half that of the Hann function and a depth equal to its peak. This significantly simplifies the estimation of the phase shift through the defect, while still providing a reasonably accurate estimate. The maximum phase shift for this homogeneous defect could then be calculated as




Figure 7. Weighted RMS average error as a percentage of plate thickness, for reconstructions from type 1 data (i.e. experimental data from a full elastic guided wave model). Errors are given for defects of (a) 10%, (b) 30% and (c) 60% of the plate thickness. A threshold for an error of 10% of the plate thickness is marked on each graph.

in [37] as $4\pi n_\delta a/\lambda$, where $n_\delta = |c_0/c - 1|$ is the difference in the index of refraction relative to the background and a is the radius; curve A marks the line where a phase shift of around π would be caused by the defect, confirming that this criterion defines well the limit of the applicability of the Born approximation.

The most notable difference between Figs. 8 and 9 is that the minimum size of the defects which can be reconstructed by HARBUT and DT is reduced when using experimental data obtained from the more realistic elastic guided wave model, reflecting a reduction in the resolution. While the actual limit varies to an extent with defect depth, it is around the $\lambda/2$ line marked B in Fig. 8 for the HARBUT reconstruction from type 2 data, while for the same reconstruction from type 1 data the limit, shown as C, is around 1.8λ . This suggests a loss of resolution in the reconstructions when using a purely acoustic inversion approach to reconstruct from the realistic elastic guided wave type 1 data, compared to using the same acoustic inversion methods with the type 2 data obtained from the approximate acoustic forward model.




Figure 8. Regions in defect size and depth space where acceptable reconstructions can be generated from HARBUT (blue), bent ray (green) and DT (red), using type 2 data (i.e. data from an acoustic approximation of the guided wave experiment).




Figure 9. Regions in defect size and depth space where acceptable reconstructions can be obtained from HARBUT (blue), bent ray (green) and DT (red), using type 1 data (i.e. experimental data from a full elastic guided wave model).




Figure 10. Photograph of the experimental setup. Two defects are machined into the steel plate, and arcs of reflective strips are placed around these. One of the transducers is then used to excite the guided wave, and the laser vibrometers at 45 degrees to the plate surface are used to measure the surface velocity of the scattered field.

(e) Experimental validation

Two experimental validation tests were performed to confirm that the 3D finite element model of the plate was producing accurate results. The configuration was as described in Sec. (a), to match the FE models, and the two defects were 120mm (3.2λ) wide and 6mm (60%) deep, and 60mm (1.6λ) wide and 3mm (30%) deep, matching Fig. 5(d) and (h). Both defects were machined in the same $1.2 \times 1.2\text{m}^2$ plate, separated by 0.5m, with interference between them unlikely to be significant in the reconstructions, and both were positioned sufficiently far from the plate edge to enable the initial wavepacket to be separated in time from any boundary reflections.

As in the simulations, due to symmetry, only a single source was needed. For this, a piezoelectric transducer was used to excite the A_0 wave; this transducer is as described in [19]. This is around 10mm in diameter (i.e. small relative to the wavelength) and predominantly excites an out-of-plane force on the upper surface. It has been shown to produce very pure A_0 guided waves for a 10mm thick plate. This transducer was bonded to the plate with epoxy to give good coupling, then excited with the toneburst directly from a Handyscope HS3 (TiePie Engineering, Sneek, The Netherlands).

Measurements were taken via a pair of Polytec OFV-505 laser Doppler vibrometers as described in [18]; these were positioned with their beams at 45 degrees to the plate surface, as illustrated in Fig. 10(a), configured to measure the out-of-plane and one component of the in-plane velocities of the plate surface. For the reconstructions, only the out-of-plane components, corresponding to the A_0 mode, were needed, so the in-plane values could be discarded for this study; while the out-of-plane components can be measured with a single laser beam normal to the surface, in this case the in-plane components were simultaneously acquired for other research. A semi-circular reflective strip was attached to the plate to maximise the light reflected back to the laser vibrometers along the scanned path and improve the signal to noise ratio. The laser heads were scanned to measure a 180 degree arc from the source transducer to a point directly opposite, with measurements taken at 65 points, corresponding to an effective 128 transducer array; this path is marked by the reflective strips visible in Fig. 10(b). Clearly, the early measurement locations coincide with the source transducer, so these terms are discarded from the measurements.




Figure 11. Processing of experimental signals for the larger experimental defect (3.2λ wide and 6mm, 60%, deep). (a) shows the 40th of 64 out-of-plane time trace (i.e. positioned at (68.9, -166.3) mm on the plate surface, while the defect is centred at the origin and the source is at (-180,0) mm) measured by the laser setup. The gating function is marked, used to remove any reflections from the boundary present in the signal. (b) Shows the time picking from the full set of 64 signals. These have had the dispersion removed as described in the text, and the arrival times are manually picked to correspond to an early peak in the signal.

Figure 11(a) presents a single time trace from the measurements. The initial, direct wavepacket is clearly visible, and the later arrivals are the reflections from the boundaries of the plate. The defects and measurement locations are positioned on the plate to avoid any reflections interfering with the signals of interest, so the signals are straightforward to separate. The traces are gated before and after to reduce the noise in the signals, as outlined in [11]. The start of the signal is approximated from the group velocity and geometric distances, neglecting the disturbances caused by the defect, and is assumed to finish 0.11ms later to capture the 5 cycles at 50kHz, extended by 10% to account for some dispersion. A cosine taper window is taken around this, smoothed via the cosine function to zero over a time of 0.04ms, equivalent to two cycles. An FFT (fast Fourier transform) algorithm is then applied to this signal to enable frequency domain data to be extracted.

Extracting the arrival times from data, as discussed earlier, is a complete research area in itself; in this example, since relatively few arrival times needed to be determined, manual picking was chosen. It is felt that manual picking is likely to give the most accurate results to enable a representative comparison to be made. To achieve this, firstly the majority of the wavepacket distortion due to dispersion was removed by performing a backwards propagation, by multiplying each signal by $\exp(-ik_0d)$ in the frequency domain, where d is the straight-line distance from the source to each receiver. If there were no defects, the arrival times of these signals should then all be zero. The arrival times are then manually picked, as illustrated in Fig. 11(b); in this case they have been chosen to follow an early peak in the signals (note that the offset to zero is subsequently removed by subtracting a constant). Finally, to account for the backwards propagation step performed initially, the arrival times must be re-propagated by adding d/c_{gp} . As with the numerical models, full matrices of frequency domain and arrival time data were produced by replicating the signals, exploiting the symmetry of the configuration.

Figure 12 presents the reconstructed images for the experimental reconstructions compared to the equivalent reconstructions from FE data. Both the HARBUT and DT reconstructions match very well between the FE model and the experimental data, although there is notable noise in the background of the experimental images. As observed in [11], in most normal scenarios any noise present in each trace will be incoherent between the different measurements in the array; when using a reconstruction algorithm, the noise in the final image will be reduced to an extent by averaging. However, the act of copying measurements to synthesise a full array of measurements




Figure 12. Comparison of imaging reconstructions (ray tomography, diffraction tomography and HARBUT) from FE and experimental data for two defects. The reconstructions from simulated data are shown on the left, in (a) and (c); these are repeated from Figs. 5(d) and (h) respectively. The reconstructions from the experimental data are shown in (b) and (d). The top row, (a) and (b) are for a defect 120mm (3.2λ) wide and 6mm (60%) deep, and the bottom row (c) and (d) are the 60mm (1.6λ) wide and 3mm (30%) deep defect.

introduces coherence between the noise components, which will not be averaged out to the same extent. Thus one would expect the effect of noise to be lower in practice than demonstrated here.

Ray tomography shows more significant differences between the simulation and the experimental data, primarily due to the difficulties in determining the arrival times of the signals, as discussed above. However, the resolution limitations of the algorithm due to diffraction are clearly visible in both sets of reconstructions, and the overall behaviour between the simulations and experiments is very similar.

4. Complex defect comparison

Previous analysis has focused on simple defects, enabling a general analysis of the performance of the algorithms to be identified. Such defects are not representative of the results in practice, however, since defects are typically complex with many features. One approach could be to consider that the complex defects are a particular combination of the simpler defects, with the resulting image being a linear combination of the images of these simple defects. However, this relies on the linearity of both the wave behaviour and the algorithm itself, which as discussed earlier cannot be guaranteed under general circumstances. Given that the assumptions in the algorithms do allow the wave scattering problem to be explained with a degree of linearity,




Figure 13. Reconstructions from the complex defect shown in (a). (b) is ray tomography, (c) is DT and (d) is HARBUT, all from type 1 data (experimental data produced from the realistic elastic guided wave simulation). (e) is the HARBUT reconstruction, instead using type 2 data (experimental data obtained from the acoustic approximation of the wave behaviour). The cross sections along the line marked in (a) are shown in (f) for the original case, the HARBUT reconstructions from type 1 and type 2 data.

however, it could be expected that the trends identified for the simple defects would be present for a complex defect too.

To investigate this problem, a realistic complex defect is calculated from a laser scanning of the surface profile of a corrosion defect found at a pipe support. A thickness map based on these measurements is presented in Fig. 13(a). At present, RAM limitations on the GPU restrict the maximum size of the model, limiting the usable model size to around $400 \times 400 \text{ mm}^2$. Therefore it is not possible to model large arrays, so this has been set to 360mm in diameter. The defect has been scaled from the original by a factor of around 0.7 in the in-plane directions to enable it to fit within this array. To achieve a maximum 50% wall loss (i.e. 5mm for the 10mm plate), the defect is similarly scaled by a factor of 0.7 in the out-of-plane direction. This defect was selected because it has a number of interesting features, including a particularly deep point close to the edge; it should therefore form a significant challenge to the reconstruction algorithms. Simulations of this model are performed using the elastic FE model and the acoustic FD equivalent, to produce type 1 and type 2 experimental data respectively. These data sets are subsequently used to test the acoustic imaging algorithms as before.

Figures 13(b), (c) and (d) present the reconstructions for ray tomography, DT and HARBUT respectively, all when using type 1 data. The ray tomography reconstruction is poor, with the smaller features present having little in common with the original thickness map, although the overall outline of the object is fairly well matched. Diffraction tomography is very poor since the contrast and the size are too large for the Born approximation to be valid. HARBUT gives a reasonable reconstruction, although it is clear that the resolution is limited, preventing the algorithm from being able to pick out the finer details, so importantly the algorithm is unable to determine the deepest point in the thickness map. For comparison, the HARBUT reconstruction from type 2 data is presented in Fig. 13(e), then cross sections along the line marked in Fig. 13(a)




Figure 14. Reconstruction from a section of the complex defect of Fig. 13(a), containing the deepest point, scaled up by a factor of four in the two in plane dimensions. (a) shows the true thickness map. (b) gives the HARBUT reconstruction, produced from type 1 data (experimental data produced from the realistic elastic guided wave simulation). (c) plots the cross sections of the original and the HARBUT reconstruction along the white dashed line marked in (a).

for the thickness maps of Figs. 13(a), (d) and (e) are plotted in Fig. 13(f). It is clear that when using HARBUT with type 2 data, the reconstruction is much higher resolution than with type 1 data, confirming the results seen with the simple defects.

To confirm that the inability to reconstruct the deepest point is down to resolution limitations when using acoustic imaging algorithms with type 1 data, the same defect is scaled up by a factor of four in each of the in-plane dimensions and a HARBUT image is generated from this data. Since this would extend beyond the edge of the 360mm diameter array, a circular window is applied to limit the extent. The resulting thickness map used is shown in Fig. 14(a). The reconstruction is presented in Fig. 14(b), and the cross sections are compared in Fig. 14(c), showing how HARBUT is now able to better capture the deepest point in the image, confirming that the failure of the algorithm previously was a result of the scale of the defect. This gives a strong indication that the rapid changes in thickness do cause errors in the assumption that the acoustic imaging algorithms can capture all the behaviour of the guided waves as they travel through the plate.

As discussed earlier, the standard resolution limit of HARBUT is $\lambda/2$, based on the Born approximation [37]. This resolution limit is defined as the period of the highest frequency component which will exist in the image, and in the case of the Born approximation the limit occurs because the reconstruction is, in effect, filtered to remove any wavenumber components above $2k_0$, where k_0 is the wavenumber of the background medium. The resolution can be determined by comparing the various reconstructions to the original in the Fourier domain. Figure 15 plots the 2D Fourier transforms of Fig. 13(a), (d) and (e) to enable the resolution to be compared.

The original, in Fig. 15(a) clearly has significant features at a wide variety of wavenumbers. Both of the HARBUT reconstructions, using type 1 experimental data (extracted from the full elastic guided wave model) in Fig. 15(b) and type 2 data (from the approximate acoustic model) in Fig. 15(c) are noticeably filtered. The approximate limits in k_x where the amplitudes of the features drop significantly below those of the original are marked by the dashed white lines. In the case of Fig. 15(c), this filtering matches the $2k_0$ limit for the Born approximation methods, which is expected since both the forward model and the inversion are performed using an acoustic model, but the limit of Fig. 15(b) is much lower. The radius where the component amplitudes drop significantly below the original is around $0.6k_0$, which suggests a resolution limit of around $1.5 - 2\lambda$, matching closely the observations made earlier in the paper for the simple defects.

5. Analysis and discussion

The results presented in this paper for a variety of different scenarios have shown that resolution is significantly limited when reconstructing from type 1 experimental data obtained from the




Figure 15. K-space (i.e. spatial frequency domain) representations of the images. These are normalised relative to the peak, and the colour scale is set to the range 0-20% to enable the small components to be seen. (a) is the original, corresponding to the FFT of Fig. 13(a). (b) gives the frequency spectrum of the HARBUT reconstruction from type 1 data (Fig. 13(d)). (c) is the frequency spectrum of the HARBUT reconstruction from type 2 data (Fig. 13(e)). The white dashed lines in (b) and (c) mark the k_x values at which the amplitudes of the features drop significantly below those of the original in (a).




Figure 16. Comparison of scattered field for a small, shallow defect (flat bottomed hole, 5% deep, $\lambda/4$ diameter) with data taken from the elastic model (blue, corresponding to type 1 data) and the acoustic equivalent (red, corresponding to type 2 data). (a) defines the scattering angle θ . (b) plots the amplitude and (c) plots the phase as a function of θ . Marked in grey is the ‘matching region’, where the amplitude of the elastic field is within 6dB of the acoustic field, and the phases are within around 0.7 radians.

realistic elastic guided wave model, when compared to the equivalent reconstruction using type 2 data extracted from the approximate acoustic forward model. To analyse this further, we can consider a small, point-like flat-bottomed defect which lies within the Born approximation. The chosen defect is 5% of the nominal thickness (0.5mm) deep and of diameter $\lambda/4$ (10mm). As before, type 1 and type 2 experimental data are produced from the full elastic guided wave model and the acoustic approximation respectively, then the resulting scattered field is calculated by subtracting the incident field. This scattered field, output for different scattering angles as defined in Fig. 16(a), is plotted in Figs. 16(b) and (c), for amplitude and phase respectively. In these, there is a clear region where the amplitude and phase match within the limit of 6dB amplitude and 0.7 radians, from scattering angles of 0 to 0.64 radians, marked in grey, although the amplitudes in particular are clearly diverging. Outside this region there is significant difference between the acoustic and the elastic models. Based on this, there is a region of transmitted measurements of around ± 0.64 radians either side of through-transmission where the acoustic model forms a reasonable approximation of the elastic data.

Under the Born approximation, as discussed in Sec. ii, the different measurements of the scattered field can be directly mapped to the 'K-space', the 2D Fourier transform of the image space plotted in Fig. 15. For a given illumination direction $\mathbf{s}_0$ and measurement direction $\mathbf{s}$ the measured scattered field component maps to the point in K-space at $k_0(\hat{\mathbf{s}}_0 - \hat{\mathbf{s}})$; the approach of diffraction tomography (and by extension, HARBUT) is to use these measurements to populate the K-space, then apply an inverse Fourier transform to obtain the reconstruction of the object. The resolution limit will be given by the highest wavenumber component present in the K-space. Given that a circular array is used to give a full view configuration, we do not expect any directional bias in the maximum wavenumber that can be obtained, so just the magnitude of the vector is considered

$$k_0 |\hat{\mathbf{s}}_0 - \hat{\mathbf{s}}| = k_0 \sqrt{(\hat{\mathbf{s}}_0 - \hat{\mathbf{s}}) \cdot (\hat{\mathbf{s}}_0 - \hat{\mathbf{s}})} \quad (5.1)$$

where $\cdot$ represents the standard dot product. Therefore

$$k_0 |\hat{\mathbf{s}}_0 - \hat{\mathbf{s}}| = k_0 \sqrt{\hat{\mathbf{s}}_0 \cdot \hat{\mathbf{s}}_0 - 2\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}_0 + \hat{\mathbf{s}} \cdot \hat{\mathbf{s}}} \quad (5.2)$$

$$= k_0 \sqrt{2 - 2 \cos \theta} \quad (5.3)$$

$$= 2k_0 \sqrt{\frac{1 - \cos \theta}{2}} \quad (5.4)$$

$$= 2k_0 \sin \frac{\theta}{2} \quad (5.5)$$

where θ is the angle between the vectors $\hat{\mathbf{s}}$ and $\hat{\mathbf{s}}_0$, and corresponds to the scattering angle marked in Fig. 16(a). From this, it is clear that the transmitted components (θ tends to zero) contain information about the low wavenumber components of the object, while the reflected components (θ tends to π) contain the highest wavenumber components, around $2k_0$. In this example, scattered data from the acoustic and elastic models only matches up to an angle of ± 0.64 radians either side of through-transmission, so all the wavenumber components within a circle of radius $2 \sin(0.64/2) = 0.63k_0$ can be obtained, which matches the plots from Fig. 15(b).

While this analysis has been undertaken under the Born approximation, the results do generalise outside it. HARBUT produces its reconstruction by using the Born approximation with corrections to allow it to work relative to a known background, and it would therefore be anticipated that the HARBUT reconstruction would be similarly limited. This is why the resolution limitation of HARBUT when using data from the elastic forward model matches this theory so well.

The practical implication of the resolution limit when using acoustic imaging algorithms for this problem is that the finer details of the thickness map will be removed. As demonstrated for the complex defect of Fig. 13 this can prevent the minimum remnant wall thickness from being accurately identified under certain circumstances. There is therefore a need to improve the resolution. One option is to use a higher frequency, to reduce the wavelength, so the resolution will improve due to the shorter wavelength. However, there are a number of practical issues associated with changing the operating point. Exciting a pure mode will become more challenging due to the presence of additional modes at higher frequencies, making transduction significantly more critical. Additionally, the sensitivity of the guided waves to thickness changes will often vary with frequency, and there can be additional issues including attenuation; a full study of these is outlined in [48]. A final point is that a shorter wavelength will require more transducers to enable the field to be sufficiently sampled.

The other potential solution for improved resolution would be to use a reconstruction approach which is more faithful to the full behaviour of the 3D guided waves, rather than relying on the limited accuracy of the guided wave to acoustic mapping. This could be achieved through a 'full wave inversion' approach, where the 3D FE model is repeatedly run and updated to match the experimental data. Clearly this is susceptible to any local minima present in the solution, and is likely to be sensitive to noise, although the issues associated with the former in particular could be minimised by using the lower resolution reconstruction from HARBUT as a starting point. It

should also be recognised that while the $\lambda/2$ resolution has been achieved when using acoustic models in both simulation of experimental data and the subsequent imaging techniques, there is no guarantee that such a resolution can be achieved when a full elastic inversion approach is used with experimental data from an elastic model.

6. Conclusions

This paper has clearly defined the performance of the ray tomography, diffraction tomography and HARBUT algorithms for thickness mapping via guided wave tomography. When using simulated experimental data generated by an acoustic model (type 2 data), it was clear that ray tomography is limited by the presence of diffraction that occurs at small scales, preventing small scale objects from being reconstructed. With the same data, diffraction tomography is limited by Born approximation, which prevents it from generating accurate reconstructions when the contrast and size combine to cause a large phase shift through the defect. HARBUT shows a good step forward over both of these, performing as well as either within their respective limits, but also producing good reconstructions in the intermediate region where neither performs well, achieving a resolution of around $\lambda/2$ across all tested defect depths. When the simulated experimental data is generated by a more realistic elastic model instead (i.e. type 1 data was used), this paper has shown that the achievable resolution drops to around $1.5\text{--}2\lambda$. The accuracy of the data produced from this elastic model was also confirmed with a set of experiments. For the defects across the ranges tested (5–60% wall loss and $0.125\text{--}9.6\lambda$ diameter), HARBUT was shown to reconstruct the peak wall loss within 10% of nominal plate thickness and have an average shape error of less than 5%, down to the resolution limit of $1.5\text{--}2\lambda$ for type 1 data, and $\lambda/2$ for type 2 data.

A physical investigation of the resolution loss when reconstructing from simulated experimental data from the elastic model (type 1 data) compared to simulated experimental data from the acoustic model (type 2) has explained how the scattered field differs between the two models used, showing that there is a limited range of angles either side of through-transmission where the difference between the two is relatively low. Analysis has shown how this range of angles only allows components up to the $1.5\text{--}2\lambda$ resolution limit to be extracted when acoustic inversion approaches are used, indicating that this is a fundamental limit which restricts the resolution of this widely used approach in guided wave tomography. A clear opportunity for future resolution improvements is to develop inversion approaches which can more accurately capture the physics of the elastic guided wave scattering occurring in the type 1 data, rather than approximating it with an acoustic model.

Acknowledgements

The author is grateful to Professor P. Cawley for the critical reading of the paper and the helpful discussions, and also to J. Isla for the assistance with the experimental measurements.

References

1. H. Carr and E. Purcell, "Effects of diffusion on free precession in nuclear magnetic resonance experiments," *Physical Review*, vol. 94, no. 3, p. 630, 1954.
2. G. Hounsfield, "Computerized transverse axial scanning (tomography): Part 1. description of system," *British Journal of Radiology*, vol. 46, no. 552, p. 1016, 1973.
3. R. P. Bording, A. Gersztenkorn, L. R. Lines, J. A. Scales, and S. Treitel, "Applications of seismic travel-time tomography," *Geophysical Journal of the Royal Astronomical Society*, vol. 90, no. 2, pp. 285–303, 1987.
4. J. K. Cohen and N. Bleistein, "An inverse method for determining small variations in propagation speed," *SIAM Journal on Applied Mathematics*, vol. 32, no. 4, pp. 784–799, 1977.
5. S. Raz, "Direct reconstruction of velocity and density profiles from scattered field data," *Geophysics*, vol. 46, no. 6, pp. 832–836, 1981.

6. L. Lines and S. Treitel, "Tutorial: A review of least-squares inversion and its application to geophysical problems," *Geophysical prospecting*, vol. 32, no. 2, pp. 159–186, 1984.
7. P. Mora, "Nonlinear two-dimensional elastic inversion of multioffset seismic data," *Geophysics*, vol. 52, no. 9, pp. 1211–1228, 1987.
8. M. Cheney, D. Isaacson, and J. C. Newell, "Electrical impedance tomography," *SIAM review*, vol. 41, no. 1, pp. 85–101, 1999.
9. D. B. Avdeev, "Three-dimensional electromagnetic modelling and inversion from theory to application," *Surveys in Geophysics*, vol. 26, no. 6, pp. 767–799, 2005.
10. A. Volker and J. Bloom, "Experimental results of guided wave travel time tomography," in *AIP Conference Proceedings*, vol. 1335, p. 215, 2011.
11. P. Huthwaite and F. Simonetti, "High-resolution guided wave tomography," *Wave Motion*, vol. 50, no. 5, pp. 979–993, 2013.
12. J. Wong, P. Hurley, and G. West, "Crosshole seismology and seismic imaging in crystalline rocks," *Geophysical Research Letters*, vol. 10, no. 8, pp. 686–689, 1983.
13. M. Worthington, "An introduction to geophysical tomography," *First Break*, vol. 2, no. 11, 1984.
14. N. Bregman, R. Bailey, and C. Chapman, "Crosshole seismic tomography," *Geophysics*, vol. 54, no. 2, pp. 200–215, 1989.
15. J. Zhang and M. N. Toksöz, "Nonlinear refraction traveltimes tomography," *Geophysics*, vol. 63, no. 5, pp. 1726–1737, 1998.
16. P. Williamson, "A guide to the limits of resolution imposed by scattering in ray tomography," *Geophysics*, vol. 56, no. 2, pp. 202–207, 1991.
17. K. Leonard and M. Hinders, "Multi-mode lamb wave tomography with arrival time sorting," *The Journal of the Acoustical Society of America*, vol. 117, p. 2028, 2005.
18. P. Belanger and P. Cawley, "Feasibility of low frequency straight-ray guided wave tomography," *NDT & E International*, vol. 42, no. 2, pp. 113–119, 2009.
19. P. Belanger, P. Cawley, and F. Simonetti, "Guided wave diffraction tomography within the born approximation," *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, vol. 57, no. 6, pp. 1405–1418, 2010.
20. P. Belanger, *Feasibility of thickness mapping using ultrasonic guided waves*. PhD thesis, PhD thesis, Imperial College, London, 2009.
21. P. Huthwaite and F. Simonetti, "High-resolution imaging without iteration: a fast and robust method for breast ultrasound tomography," *Journal of the Acoustical Society of America*, vol. 130, no. 3, pp. 1721–1734, 2011.
22. P. Huthwaite, F. Simonetti, and N. Duric, "Combining time of flight and diffraction tomography for high resolution breast imaging: Initial invivo results (I)," *The Journal of the Acoustical Society of America*, vol. 132, p. 1249, 2012.
23. H. Lamb, "On waves in an elastic plate," *Proceedings of the Royal Society of London. Series A*, vol. 93, no. 648, pp. 114–128, 1917.
24. M. J. S. Lowe, "Matrix techniques for modeling ultrasonic waves in multilayered media," *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, vol. 42, no. 4, pp. 525–542, 1995.
25. B. Pavlakovic, M. J. S. Lowe, D. Alleyne, and P. Cawley, "Disperse: a general purpose program for creating dispersion curves," *Review of progress in quantitative nondestructive evaluation*, vol. 16, pp. 185–192, 1997.
26. R. Pratt, "Seismic waveform inversion in the frequency domain, part 1: Theory and verification in a physical scale model," *Geophysics*, vol. 64, no. 3, pp. 888–901, 1999.
27. V. Burov, N. Alekseenko, and O. Rumyantseva, "Multifrequency generalization of the novikov algorithm for the two-dimensional inverse scattering problem," *Acoustical Physics*, vol. 55, no. 6, pp. 843–856, 2009.
28. S. Fomel, "Traveltimes computation with the linearized eikonal equation." http://sepwww.stanford.edu/public/docs/sep94/sergey3/paper_html/index.html.
29. J. Sethian, "A fast marching level set method for monotonically advancing fronts," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 93, no. 4, p. 1591, 1996.
30. J. Sethian, "Fast marching methods," *SIAM review*, pp. 199–235, 1999.
31. P. S. Earle and P. M. Shearer, "Characterization of global seismograms using an automatic-picking algorithm," *Bulletin of the Seismological Society of America*, vol. 84, no. 2, pp. 366–376, 1994.
32. D. A. Dodge, G. C. Beroza, and W. Ellsworth, "Foreshock sequence of the 1992 landers, california, earthquake and its implications for earthquake nucleation," *Journal of Geophysical Research: Solid Earth (1978–2012)*, vol. 100, no. B6, pp. 9865–9880, 1995.

33. R. Sleeman and T. van Eck, "Robust automatic p-phase picking: an on-line implementation in the analysis of broadband seismogram recordings," *Physics of the earth and planetary interiors*, vol. 113, no. 1, pp. 265–275, 1999.
34. H. Zhang, C. Thurber, and C. Rowe, "Automatic p-wave arrival detection and picking with multiscale wavelet analysis for single-component recordings," *Bulletin of the Seismological Society of America*, vol. 93, no. 5, pp. 1904–1912, 2003.
35. S. Caorsi, A. Massa, and M. Pastorino, "A computational technique based on a real-coded genetic algorithm for microwave imaging purposes," *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 38, no. 4, pp. 1697–1708, 2000.
36. X. Chen, D. Liang, and K. Huang, "Microwave imaging 3-d buried objects using parallel genetic algorithm combined with fdtd technique," *Journal of Electromagnetic Waves and Applications*, vol. 20, no. 13, pp. 1761–1774, 2006.
37. A. Kak and M. Slaney, *Principles of computerized tomographic reconstruction*. New York: IEEE Press, 1998.
38. S. Li, K. Mueller, M. Jackowski, D. Dione, and L. Staib, "Fast marching method to correct for refraction in ultrasound computed tomography," in *3rd IEEE International Symposium on Biomedical Imaging: Nano to Macro, 2006*, pp. 896–899, 2006.
39. A. Hormati, I. Jovanovic, O. Roy, and M. Vetterli, "Robust ultrasound travel-time tomography using the bent ray model," in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, vol. 7629, p. 17, 2010.
40. P. R. Williamson and M. Worthington, "Resolution limits in ray tomography due to wave behavior: Numerical experiments," *Geophysics*, vol. 58, no. 5, pp. 727–735, 1993.
41. G. Nolet and F. Dahlen, "Wave front healing and the evolution of seismic delay times," *Journal of geophysical research*, vol. 105, no. B8, pp. 19043–19, 2000.
42. M. Born and E. Wolf, *Principles of Optics*. Cambridge: Cambridge University Press, 1999.
43. J. Goodman, *Introduction to Fourier Optics*. New York: McGraw-Hill, 1996.
44. P. Huthwaite, "Accelerated finite element elastodynamic simulations using the GPU," *Journal of Computational Physics*, vol. 257, Part A, no. 0, pp. 687 – 707, 2014.
45. M. Drozd, *Efficient Finite Element Modelling Of Ultrasound Waves in Elastic Media*. PhD thesis, Imperial College London, 2008.
46. P. Rajagopal, M. Drozd, E. A. Skelton, M. J. Lowe, and R. V. Craster, "On the use of absorbing layers to simulate the propagation of elastic waves in unbounded isotropic media using commercially available finite element packages," *NDT & E International*, vol. 51, pp. 30–40, 2012.
47. D. Komatitsch and R. Martin, "An unsplit convolutional perfectly matched layer improved at grazing incidence for the seismic wave equation," *Geophysics*, vol. 72, no. 5, pp. 155–167, 2007.
48. P. Huthwaite, R. Ribichini, P. Cawley, and M. J. S. Lowe, "Mode selection for corrosion detection in pipes and vessels via guided wave tomography," *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, vol. 60, no. 6, pp. 1165–1177, 2013.