Investigation of Synthetic Aperture Methods in Ultrasound Surface Imaging Using Elementary Surface Types

W. Kerr[a], S.G. Pierce[b], P. Rowe[a]

[a]*Department of Biomedical Engineering, University of Strathclyde, Glasgow, UK*
[b]*Department of Electronic and Electrical Engineering, University of Strathclyde, Glasgow, UK*

Abstract

Synthetic aperture imaging methods have been employed widely in recent research in non-destructive testing (NDT), but uptake has been more limited in medical ultrasound imaging. Typically offering superior focussing power over more traditional phased array methods, these techniques have been employed in NDT applications to locate and characterise small defects within large samples, but have rarely been used to image surfaces. A desire to ultimately employ ultrasonic surface imaging for bone surface geometry measurement prior to surgical intervention motivates this research, and results are presented for initial laboratory trials of a surface reconstruction technique based on global thresholding of ultrasonic 3D point cloud data. In this study, representative geometry artefacts were imaged in the laboratory using two synthetic aperture techniques; the Total Focusing Method (TFM) and the Synthetic Aperture Focusing Technique (SAFT) employing full and narrow synthetic apertures, respectively.

Three high precision metallic samples of known geometries (cuboid, sphere and cylinder) which featured a range of elementary surface primitives were imaged using a 5MHz, 128 element 1D phased array employing both SAFT and TFM approaches. The array was manipulated around the samples using a precision robotic positioning system, allowing for repeatable ultrasound derived 3D surface point clouds to be created. A global thresholding technique was then developed that allowed the extraction of the surface profiles, and these were compared with the known geometry samples to provide a quantitative measure of error of 3D surface reconstruction. The mean errors achieved with optimised SAFT imaging for the cuboidal, spherical and cylindrical samples were 1.3 mm, 2.9 mm and 2.0 mm respectively, while those for TFM imaging were 3.7 mm, 3.0 mm and 3.1 mm, respectively. These results were contrary to expectations given the higher information content associated with the TFM images. However, it was established that the reduced error associated with the SAFT technique was associated with significant reductions in side lobe levels of approximately 24dB in comparison to TFM imaging, although this came at the expense of reduced resolution and coverage.

Keywords: Ultrasound imaging, Total Focussing Method, Synthetic Aperture Focussing Method, Full Matrix Capture, Robotics.

1. Introduction

A number of robotically guided knee arthoplasty systems require a preoperative 3D model of the joint, with which the surgery can be planned and implemented [1]–[4]. Computed Tomography (CT) is seen as the ‘gold standard’ in this area, but it is costly [5] and can apply a dose of ionising radiation greater than the yearly background dose of 2.2 mSv [6], which is potentially dangerous to the patient [7]. Ultrasound imaging has the potential to provide an alternative to CT in this capacity by offering comparable accuracies, while reducing cost and eliminating the risk associated with ionising radiation.

Synthetic aperture imaging methods have become commonplace in research in non-destructive testing (NDT), allowing for improved focussing capability and increased resolution over more traditional B-scan methods [8]. With these attributes, such techniques could provide an improvement over traditional medical imaging methods in accurately reconstructing the bony surfaces of the knee. However, these techniques have found little uptake in medical ultrasound imaging, with standard commercial systems lacking the versatility to perform Full Matrix Capture (FMC) [9] – a requirement of popular reconstruction algorithms such as the Total Focussing Method (TFM). Additionally, real time, high resolution imaging is usually a requirement of medical ultrasound systems. This is difficult to achieve using synthetic aperture methods, in that the techniques are inherently computationally expensive. As such, high frame rates often cannot be achieved, even when exploiting the parallelisable nature of the calculations [10].

Research into synthetic aperture methods in NDT has, for the most part, concentrated on locating and characterising small defects within relatively large samples [11]–[13]. While TFM has been used for surface imaging in dual media compensation calculations [14], [15], these and other auto focussing techniques have been limited to relatively simple, continuous surfaces [16], [17]. Imaging complex surfaces such as bones would, on the other hand, require surface reconstruction of highly variable and often discontinuous surface types [18]. A further challenge is to fully represent the entire surface under inspection, requiring, firstly, a high number of images and, secondly, an accurate probe positioning system.

To meet these challenges and to test the ability of synthetic aperture methods to reconstruct surfaces at a fundamental level, three high precision metallic samples of known, simple geometries were imaged. These samples provided elementary surface types and features that would be found in complex samples, such as bone. Data was captured using FMC and was processed using both TFM and a form of the Synthetic Aperture Technique (SAFT). These methods provided a contrast between the two extremes of the spectrum in synthetic aperture methods – TFM comprising a fully populated transmit-receive matrix and SAFT a minimally diagonally populated transmit-receive matrix. As can be seen in Fig. 1, TFM offers the maximum possible focussing power by synthetically focussing using the full aperture in reception. SAFT methods, on the other hand, employ a sub aperture which diminishes the focussing capabilities and imposes a higher level of positional dependence on the reconstruction. High data throughput was achieved using an FPGA-based phased array controller, while high speed data processing times were made possible through a Graphics Processing Unit (GPU) implementation of the synthetic aperture algorithms [19]. The probe was manipulated using a robotic precision positioning system, which provided accurate positional data, allowing for ultrasound derived 3D surface point cloud reconstruction. The 3D reconstructed surfaces were compared to the known reference models and the performances of the imaging methods were compared.

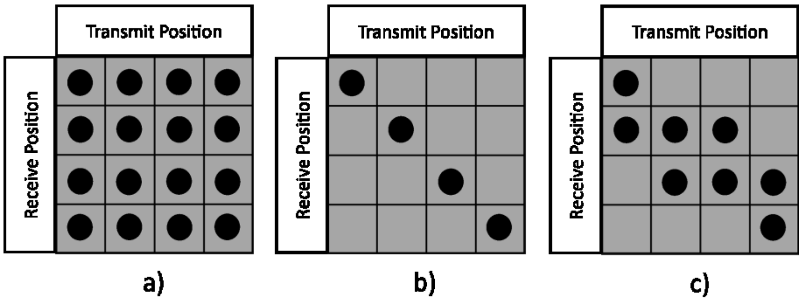

Fig. 1: Transmit-receive matrices of a four element array for TFM (a), single element receive SAFT (b) and multiple element receive SAFT (c).

The structure of the paper is, firstly, an introduction to the data capture methods and imaging algorithms employed, which is then followed by a description of the experimental apparatus and method. The results for each sample are then presented and discussed separately, after which a discussion of performance limitations is presented. This is followed, finally, by a summary and conclusions.

2. Materials and Methods

2.1 Synthetic Aperture Methods

The concept of FMC [8] is to excite one element of the phased array and receive on all the others. The succeeding element is then fired and all elements become receivers once again. This process is repeated for all N elements, producing a $N \times N$ matrix of time signals, which is known as the full matrix.

While real time synthetic aperture implementations have been achieved [10], [20], high resolution images created using probes with high element counts must still be produced in post processing. With the full matrix, it is possible to apply numerous processing methods to the same data set. One popular method is TFM, which employs every element in the array (ie. the full aperture) to synthetically focus in transmission and reception for every pixel in the image. This process begins by discretising the image region into a grid of points, each of which defines the location of a pixel in a scalar image. The intensity of a particular pixel can be calculated using Equation 1, in which $S_{i,j}$ is the time-trace associated with a transmission at the ith element and reception at the jth element, while (x, y) are the coordinates of the pixel. The time of travel from the transmitting element to the pixel is represented by $T_{i(x,y)}$, while that from the pixel to the receiving element is signified by $T_{j(x,y)}$. This is summed over the number of elements in the array, N, and is then repeated for every pixel in the image.

$$I(x, y) = \sum_{i,j=1}^{N} S_{i,j}(T_{i(x,y)} + T_{j(x,y)}) \tag{1}$$

In addition to TFM, a form of SAFT was employed in which the same process was carried out, but focussing was not performed with the full aperture. Instead, the elements constituting the synthetic aperture were defined by the position of the pixel in question. As shown in Fig. 2, while the TFM aperture included every element, the SAFT aperture was restricted to those elements contained within an isosceles triangle defined by the angle Θ.

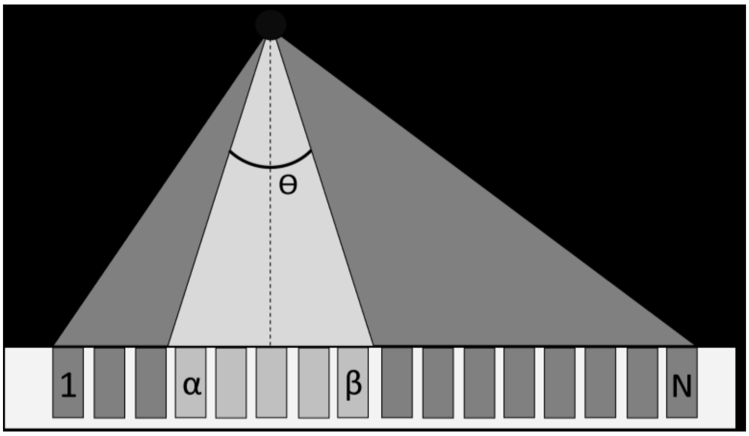

Fig. 2: Graphical representation of the TFM and SAFT synthetic aperture definitions, with the elements constituting the SAFT aperture shown in a lighter shade.

The intensity of a pixel in the image is, then, given by:

$$I(x, y) = \sum_{i=1}^{N} \sum_{j=\alpha}^{\beta} S_{i,j}(T_{i(x,y)} + T_{j(x,y)}) \tag{2}$$

where α and β are the first and last elements of the aperture, as displayed in Fig. 2. These values vary depending on the definition of Θ, with an increase in the angle increasing the aperture size.

Different SAFT aperture widths were trialled at 10° intervals from 10° to 180°, allowing for an assessment of performance and characteristics that was representative of the full spectrum of possible aperture widths. While different surface types affected performance, the 20° aperture SAFT was found to offer the greatest contrast in performance to the full aperture reconstruction of TFM whilst maintaining the ability to reconstruct most surfaces.

2.2 Experimental Apparatus

Three representative surface geometry forms were considered to encompass interaction with curves, flat surfaces and edges which are the primitive geometry forms encountered in bone geometries associated with the knee. In order to test these surface types both in isolation and in combination, three samples were prepared, the first of which was a brass sphere with a diameter of 25.0 mm (Dejay Distribution Ltd., Cornwall, UK). In addition to this, a cuboidal sample was manufactured in aluminium, with dimensions 25.1 mm × 35.4 mm × 35.2 mm. Finally, a cylindrical aluminium sample was produced with a height of 62.5 mm and a diameter of 50.0 mm. The cylindrical sample was, additionally, flattened off 15.0 mm from the centre. In doing this, the sample provided all three surface features, as can be seen in Fig 3. To ensure the majority of the surface was accessible during inspection, a mount was manufactured that allowed the samples to be elevated, such that the interrogation array probe had good line of sight access to the whole of the samples.

Fig. 3: The spherical, cuboidal and cylindrical samples shown with the mount.

A 128 element Vermon (Tours, France) phased array with a centre frequency of 5 MHz was employed for ultrasonic acquisition. This offered both a large aperture width of 89.6 mm and a theoretical resolving limit of less than 0.2 mm in water. Standard preoperative CT scans employ 1.0 mm thickness slices [1], [21]–[23], limiting the resolution to 1.0 mm. The probe was, therefore, significantly within the required spatial resolution limit, relative to reference CT imaging.

The array was excited and interrogated by a FlawInspecta (Diagnostic Sonar Ltd., Livingston, UK) phased array controller. The platform is modular, with parallel digitisation achieved using FlexRIO FPGA cards. The configuration used herein employed two 32 channel digitisers, allowing for parallel reception of 64 elements, with a maximum of 4 elements in simultaneous transmission. Therefore, two firings were required for each transmission event when using all 128 elements. This was the only limiting factor on the Pulse Repetition Frequency (PRF) originating from the hardware, with the only other constraint being that of wave travel within the material [24]. However, at the time of data capture, the firmware was not optimised, meaning that a bottleneck was created in data transfer [25] which resulted in a variable frame rate of approximately 0.3 Hz. Using a the speed of sound in water (1480 ms^{-1} [26]), images with a width of 13.42 cm and a depth of 8.95 cm were reconstructed.

In order to reconstruct 3D surfaces using 2D images, accurate probe manipulation and positional recording were vital. To this end, a KUKA KR 5 arc HW industrial robot was employed, providing six degrees of freedom and the ability to implement a range of array probe paths. Industrial robotic manipulators are seeing increasing application in high precision manipulation tasks [27]–[29] despite known issues with absolute accuracy. Despite this shortcoming, such robots are extremely repeatable in position, and have great advantages in being able to move in complex curved paths with 6 degrees of freedom whilst maintaining constant standoff and normality to the local surface geometry. Additionally the ability to employ CAD/CAM based off line programming allows ease of programming to produce complex tool paths specific to a part with known geometry [30], [31].

The correct calibration of the inspection probe tool centre point (TCP) was critical to the attained accuracy, as small errors in this physical position with respect to the flange, translated into much larger errors in absolute position as the probe was located at the end of the physical kinematic chain. A standard two stage KUKA TCP calibration method was employed [32]. The first phase of the procedure defined the position of the TCP relative to the flange and involved moving the origin of the tool reference frame to a static reference point four times, each with a different robot pose. This was achieved using a spike which was manufactured such that the position of the tip relative to the flange corresponded to the centre of the probe face. The second stage allowed for tool orientation calibration, which involved moving points on the X axis and the X-Y plane of the tool coordinate system to the reference point. Two further calibration spikes were employed to accomplish this, which conformed to the described requirements. The X-Y plane was defined such that it corresponded to the 2D imaging plane associated with the probe. Two bespoke probe holders were manufactured; one of which was parallel to the 6th axis of the robot, while the other was perpendicular, allowing for full line of sight access to the sample while maintaining full probe submersion. Each probe mount was calibrated as described above, producing calibration errors of 0.6 mm and 0.7 mm, each below the recommended industry maximum error of 0.8 mm.

To independently assess the true position of the probe compared to the reported KUKA position during manipulation, six Vicon T160 cameras (Vicon Motion Systems, Oxford, UK), were employed. A marker set of five 12 mm retro reflective markers allowed for the TCP position to be tracked while the TCP followed a hemispherical scan path with a similar working volume to that required to image the samples. The path required that a range of complex poses be adopted by the robot, which imposed highly varied joint angle combinations. The KUKA TCP position was relayed to the host PC every 12 ms via ethernet, using the Robot Sensor Interface (RSI) software [33]. The corresponding, temporally synchronised Vicon measurement was attached to each KUKA measurement. The resulting point clouds were matched using Iterative Closest Point (ICP) to compensate for the difference in the origin and orientation of the coordinate systems of the two systems. The absolute Euclidean distances between each of the corresponding KUKA and Vicon derived positions were then calculated, resulting in a mean error of 0.5 mm. This was considered a worst-case error, as the path was highly complex which would be expected to produce larger errors than would be found in more simple paths. While optical tracking is known to produce relatively large errors, these are dependent on the size of the measurement volume [34]. As the measurement and calibration volumes were small (0.002 m^2 and ~6 m^2), these errors were minimal.

In the current application, there is no a-priori information available that would allow for probe trajectory path planning for the geometry of the sample to be scanned. Indeed in the ultimate desired application in knee joint imaging, all that would be available would be the rough working volume definition around the subject's limb. Therefore, complex path programming was not required or indeed possible, and a simple rectilinear scan path around each object was employed. This had the advantage of simplification of programming using the inbuilt KUKA Robot Language (KRL) – a BASIC-like, domain specific language which allowed for simple tool paths to be defined. Such a path is shown in Fig. 4 for inspection of the cylindrical sample. The path, while not maintaining normality or a specific standoff, provided full coverage of the surface.

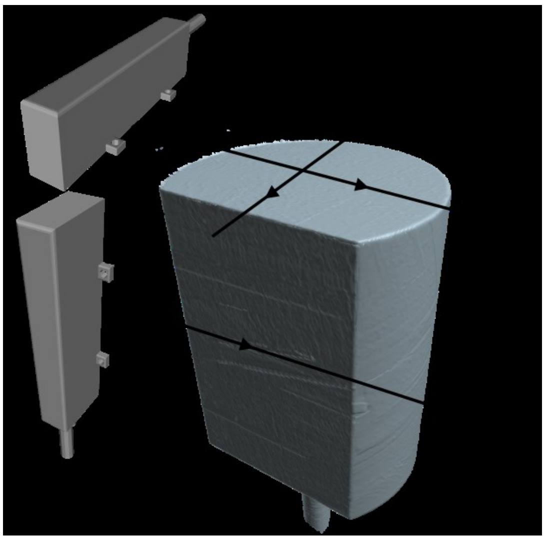

Fig. 4: Cuboidal tool path written in KRL.

The samples were placed in a water bath to allow for complete submersion, as can be seen in Fig. 5. The KRL code describing the path required initial coordinates to be defined. These were identified by manually positioning the probe to a point at which the probe face would remain submersed at all times. Additionally, the length, width and height of the cuboid defining the path were altered with each sample, so that a minimum standoff of at least 20 mm was maintained.

Fig. 5: Experimental setup, showing the submerged sample, with the probe face fully submerged.

2.3 Image and Surface Point Cloud Construction

The process of constructing images using synthetic aperture methods is computationally expensive, but is highly parallelisable [20], [35]. A software platform, cueART, has been developed by the Centre for Ultrasonic Engineering (CUE) at the University of Strathclyde, which allows for significant reductions in computation time by implementing the algorithms on a GPU [14], [19]. In doing so, cueART allowed for high resolution images to be produced using 128 element FMCs in a practical time frame. An image depth of 8

To identify the surface profile in each image, global thresholding was employed, which revealed the coordinates of all the pixels with intensities above the defined decibel limit. However, the images did not display only the first surface reflection, but multiple others, as can be seen in Fig. 6 (a). The first surface was the region of high intensity closest to the probe face (signified by circles), which represented the true location of the outer surface of the sample. The second surface, directly beneath the first, was caused by reflections from the back wall of the sample. The most prominent of the surface profiles in the image, seen at the bottom of the image, was caused by the reflected waves from the first surface reflecting on the probe face and making the return journey. As such, the line was twice the distance from the probe face as the true surface representation. In addition to these, there were numerous other false interface indications caused by further back wall reflections.

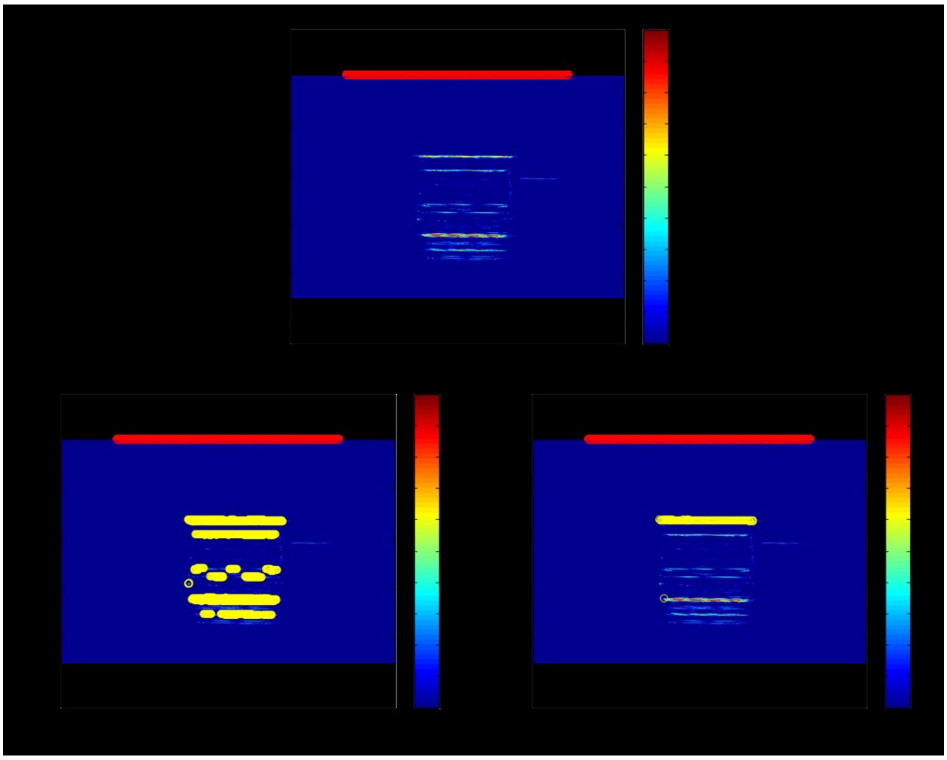

Fig. 6: A typical SAFT image obtained from the inspection of the cuboidal sample (a). Thresholding alone extracted all the erroneous surface representations present in the image (b), while the surface extraction algorithm isolated the true surface (c).

When global thresholding was performed on such an image, all the areas of high intensity were identified, as can be seen in Fig. 6 (b). In order to isolate the coordinates of the first, true surface, all the coordinates above the threshold were stored in a matrix. Any coordinates which contained the same X value were discarded, with the exception of that with the smallest Z value. As such, it was ensured that for every column of pixels, only the pixel closest to the probe face and above the decibel limit would be recorded. The result of this can be seen in Fig. 6 (c), where the erroneous surfaces have been eliminated.

The 2D coordinates representing surface contours from each image were recorded relative to the centre of the probe face. To place the points in 3D space relative to the coordinate system of the robot, the coordinates were rotated then translated using the corresponding measured KUKA position and orientation. This is presented in equation (3), where the subscripts f, o and T represent the final 3D coordinates, the original coordinates and the TCP coordinates needed for translation, respectively. Additionally, R_{xyz} is the rotation matrix in x, y, z order. Coordinate conversion from 2D to 3D was simple, as the TCP was calibrated such that the position and orientation corresponded directly to the imaging plane.

$$\begin{pmatrix} X_f \\ Y_f \\ Z_f \end{pmatrix} = R_{xyz} \begin{pmatrix} X_o \\ Y_o \\ Z_o \end{pmatrix} + \begin{pmatrix} X_T \\ Y_T \\ Z_T \end{pmatrix} \tag{3}$$

Finally, the surface point clouds were imported into CloudCompare (EDF/Telecom ParisTech, Paris, France) – an open source software package designed for comparing 3D point cloud data. The point clouds were matched with the reference models, firstly, by manually manoeuvring the cloud so that an approximate match was achieved. The point cloud was then finely matched by way of ICP. This was necessary step as the position and orientation of the samples was unknown relative to the KUKA coordinate system. While it would be advantageous to accurately position the sample such that the true position of the sample would be known relative to every packet of KUKA positional data, the samples were positioned manually and approximately to more realistically produce the final application, where no such reference data would be available from which to gain absolute positional data. It should also be said that while CloudCompare provided a convenient platform on which to compare point cloud data and visualise results, it was not vital to the system, as it did not contribute to the reconstruction process and its processes could be simply replicated on other platforms.

The error for each point in the ultrasound-derived point cloud was then calculated by finding the absolute Euclidean distance between the point and the nearest vertex on the surface of the reference model after matching. From these values, mean error, maximum error and standard deviation were calculated for each ultrasound-derived surface point cloud.

3. Results and Discussion

3.1 Cuboidal Sample

A pronounced difference can be seen in Fig. 7 between the images produced by SAFT and TFM. The SAFT-derived image shows a clear outer surface representation, along with other spurious reflections, as described in section 2.3. While the TFM image displays the same surface representations, it also exhibits side lobes of high intensity. These side lobes were artefacts generated in the reconstruction algorithm and attributed to true reflectors [19]. The typical maximum side lobe intensity for the side lobes found in the TFM images was -6 dB, while that for the SAFT images was -30 dB. The consequence of this was that when contour extraction was employed, the true surface was often not identified for both SAFT and TFM reconstructions, with the artefacts above it instead being extracted. In an effort to limit the effect of this, the threshold limit of the TFM images was set to -5 dB, while that for the SAFT images was -12 dB (the difference reflecting the side lobe levels). In doing this, the number of pixels representing the surfaces was reduces, therefore resulting in an undesirable reduction in coverage.

Additionally, it can be seen in Fig. 7 (a) that the lower surface representation appears to display a higher intensity than the upper, true surface. The reason for this is that the pixels with a larger Z coordinate would have employed a greater number of receiving elements during reconstruction, implying a greater scalar intensity. As can be seen in Fig. 2, as the pixel moves further from the probe face, the angle, Θ, remains constant, the width of the base of the triangle increases and, therefore, the number of elements in the receiving aperture increases.

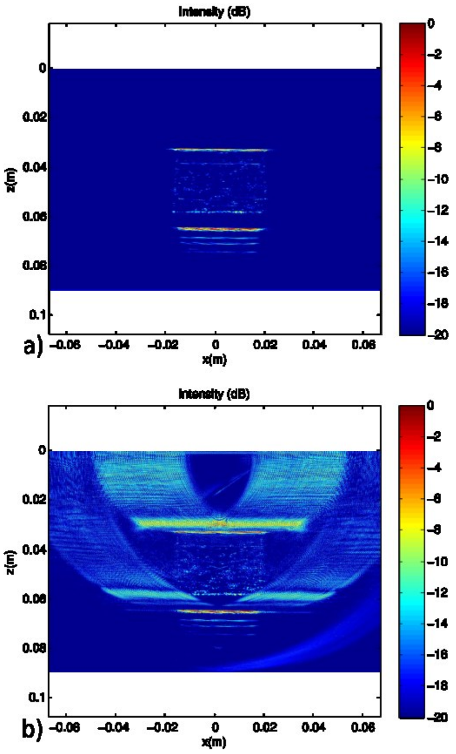

Fig. 7: SAFT (a) and TFM (b) images resulting from a typical FMC from the cuboidal sample.

Once constructed and compared, the SAFT point cloud achieved significantly lower errors than TFM, as can be seen in Table 1. Additionally, the TFM point cloud was not as dense as the SAFT counterpart, which was expected, given the higher threshold limit. Further, the TFM cloud appeared to contain multiple surfaces, as can be seen in Fig. 8 (b). The first cause of this was high intensity side lobes, while the second cause was extraction of incorrect surfaces in each image, which was made possible by the higher threshold limit associated with TFM. Comparing parts (a) and (b) of Fig. 8, it can be seen that SAFT processing provided a significantly more accurate depiction of the sample. As can be seen in Fig. 8 (a), a cuboid-like structure is present, with most of the points in these regions achieving sub-millimetre accuracy. However, there were outlying points which were a result of noise, registered due to the low threshold limit used in the SAFT images. In addition to this, each of the faces extended beyond the edge, which increased the mean error and standard deviation significantly.

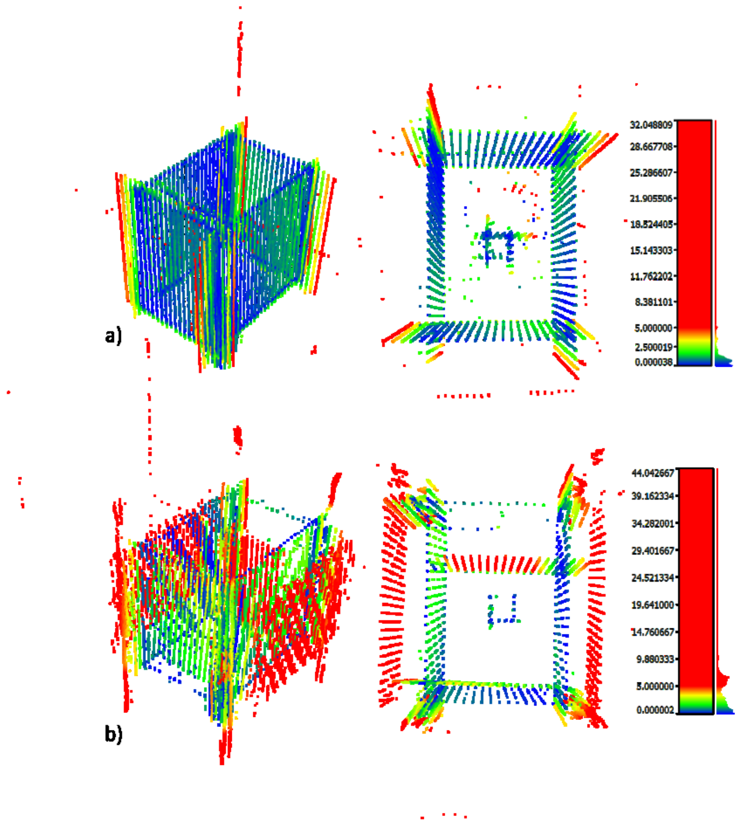

Fig. 8: SAFT (a) and TFM (b) derived surface point cloud reconstructions of the cuboidal sample, showing errors at two viewing angles.

The cause of this was not positional inaccuracy on the part of the robot controller or in the point cloud construction procedure, but rather the physical nature of the ultrasound beam itself. The beam shape of an individual element of the array has an associated thickness, not only in the width along the direction of the full aperture, but also associated with the elevation of each element. This beam width is not considered during 3D construction, as the images are regarded as 2D.

The effect of this beam width is usually negligible when imaging planar surfaces at right angles. The reason for this is that for relatively polished surfaces, such as that employed herein, specular reflection dominates. Therefore, any off-axis ultrasonic energy is directed away from the receiving element and only the energy within the imaging plane is recorded, as illustrated in Fig. 9 (a).

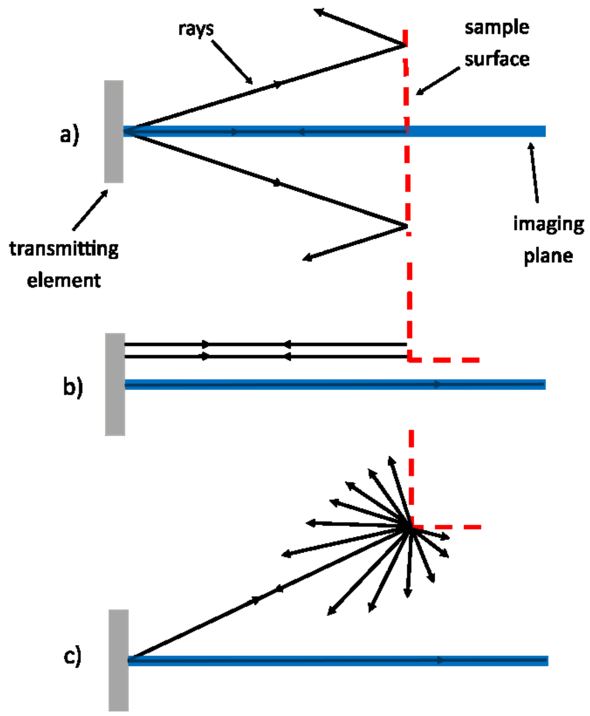

Fig. 9: Different reflection types from a planar surface (a, b) and an edge (c).

However, when the imaging plane initially passes the edge of a sample, the main beam from the element still reflects from the planar surface. As such, the resultant image will display a surface, but this will be inaccurately placed due to the assumption of a 2D imaging plane positioned through the centre of the element, as demonstrated in part (b) of Fig. 9. Given that the elevation of the elements is 10 mm, this effect would account for surfaces being registered at most 5 mm past the edge of the sample. However, in the SAFT-derived point cloud (the most accurate of the two), surfaces were reconstructed over 9 mm from the edge of the sample. This was caused by reflections from the edge of the sample, resulting from off-axis transmissions. While specular reflection dominates in planar surfaces, in the case of an edge, diffuse reflection dominates, as illustrated in Fig. 9 (c). While the intensity of the signal returned in this instance would be lower than that from the planar surfaces, the images were evaluated in a decibel scale on an individual basis, meaning the overall lower amplitude of the signals would have little impact.

3.2 Spherical Sample

The TFM images resulting from the scan of the spherical sample displayed significant side lobes, as seen in Fig. 10 (a). The most prominent of these formed as two islands of intensity similar to those that would be expected in from a point reflector. As can be seen in Fig. 10 (b), while the SAFT image showed no side lobes, there was a reduction in the resolution of the surface profile relative to the TFM image, with a thickening of the surface representation. This was a result of a reduction in the resolving power associated with SAFT relative to TFM, which offered the maximum possible resolving power [36]. Additionally, TFM allowed for reconstruction of more of the surface than SAFT. This was because the narrow synthetic aperture of SAFT only allowed for reconstruction of surfaces directly below the probe face.

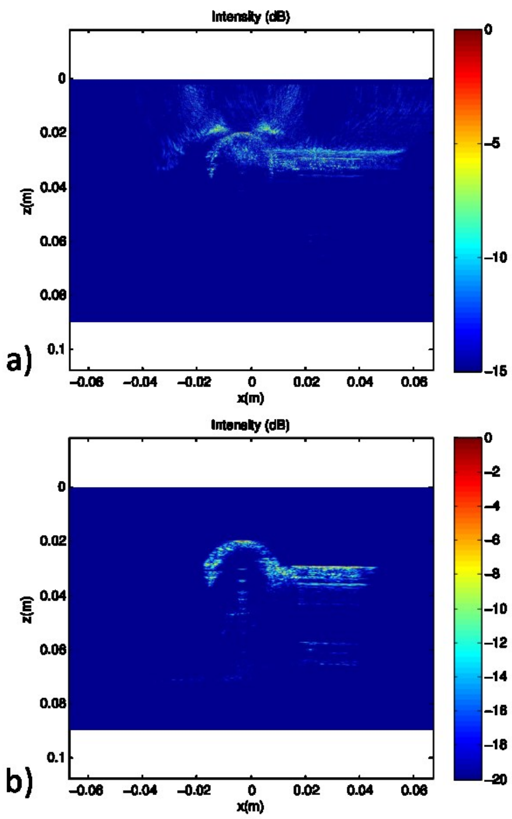

Fig. 10: TFM (a) and SAFT (b) images resulting from a typical FMC from the spherical sample.

When the probe moved away from the centre of the sphere, the obvious surface profiles exemplified in Fig. 12 were not present. Instead, the profiles became distorted, as can be seen in Fig. 11 (a). The surface extraction and 3D reconstruction of these images resulted in a lack of curvature in both the sphere and the supporting rod. This is demonstrated in Fig. 11 (b), which shows the point cloud resulting from one straight scan line of the full scan. It would be expected that the points would display an obvious curvature, but they instead possessed an almost complete lack of curvature. The same effect was encountered with the mount in the cuboidal sample, as seen in Fig. 8.

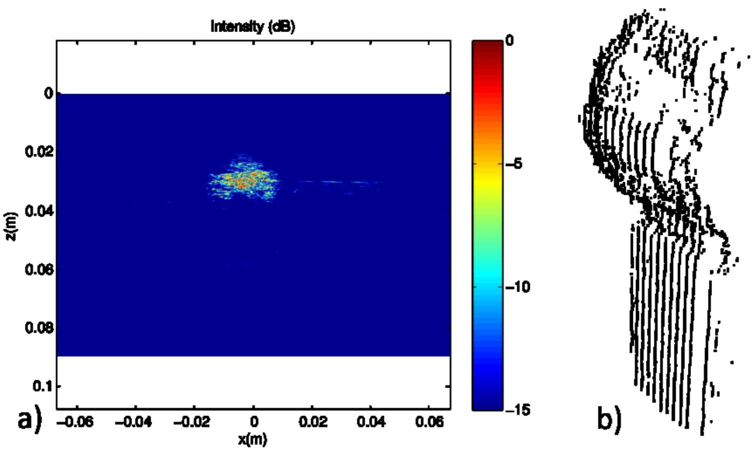

Fig. 11: SAFT image from FMC captured off the central axis of the sphere (a) and the surface reconstruction of one of the scan lines (b).

The distortion effect was caused by the three dimensional nature of the transmitting beam. As illustrated in Fig. 12, the rays radiating from the image plane were reflected away from the receiver as the probe moved away from the centre of the sphere. However, the rays outside the image plane were received and placed inaccurately. Numerous ray paths along the face of the transducer caused many reflections to be received at different times, resulting in the distortion effect. As the surface extraction algorithm discriminated in favour of those pixels closest to the probe face, the reflections originating close to the centre of the sphere were always favoured.

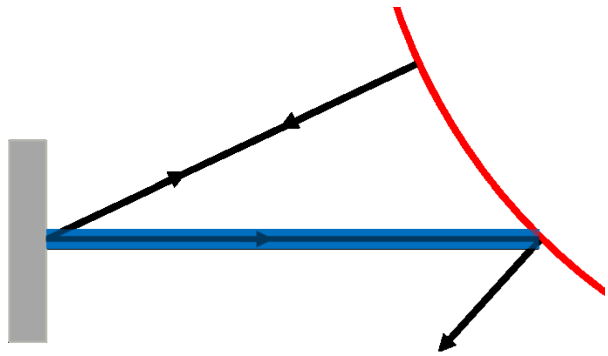

Fig. 12: Specular reflection resulting from a curved surface.

These effects had a significant impact on the reconstructed surface point clouds, with an increase in the SAFT errors relative to the cuboidal results, as shown in Table 1. While the SAFT point cloud maintained a lower mean error than the TFM cloud, it produced a significantly higher standard deviation and maximum error. It can be seen in Fig. 13 that the effects of specular reflection from the curved surface were more pronounced in the SAFT data than the TFM, causing the higher levels of error. The TFM point cloud also had a large number of inaccurate points, but they were of a different nature, being relatively close proximity to the true surface. This was, once again, caused by significant side lobes in the TFM images.

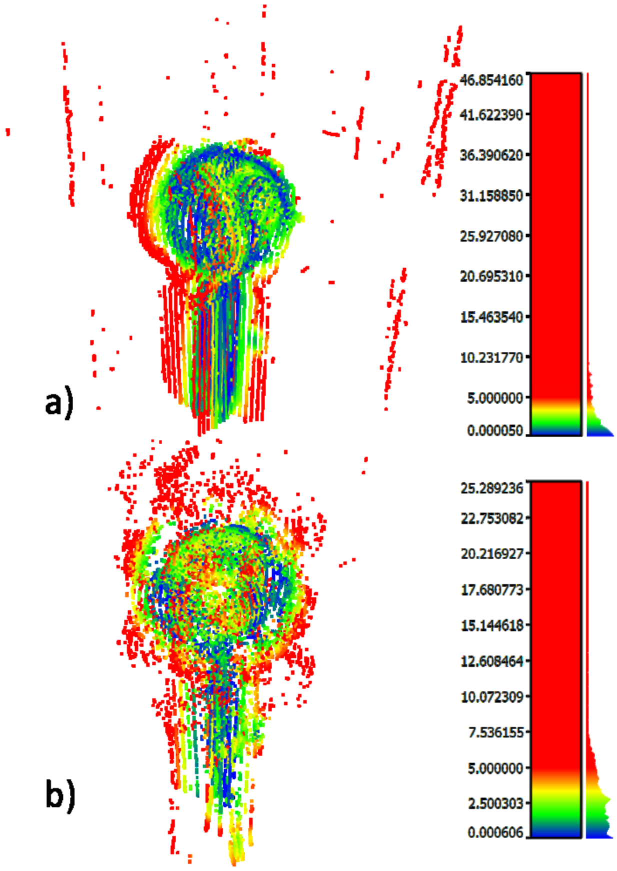

Fig. 13: SAFT (a) and TFM (b) derived surface point cloud reconstructions of the spherical sample.

3.3 Cylindrical Sample

The cylindrical sample included all three representative surface primitives (curves, flat surfaces and edges), providing the opportunity to assess them in combination. The effect of edges can be seen in Fig. 14 (a), where, towards the left of the image, inaccurate lines of points can be seen extending past the edges of the sample. The flat surfaces were generally accurate. However, it can be seen in part (a) that most of the flat surface at the top of the sample had errors of approximately 1 mm. This was most likely caused by the matching procedure, which minimises the error for all points, rather than merely those of high accuracy.

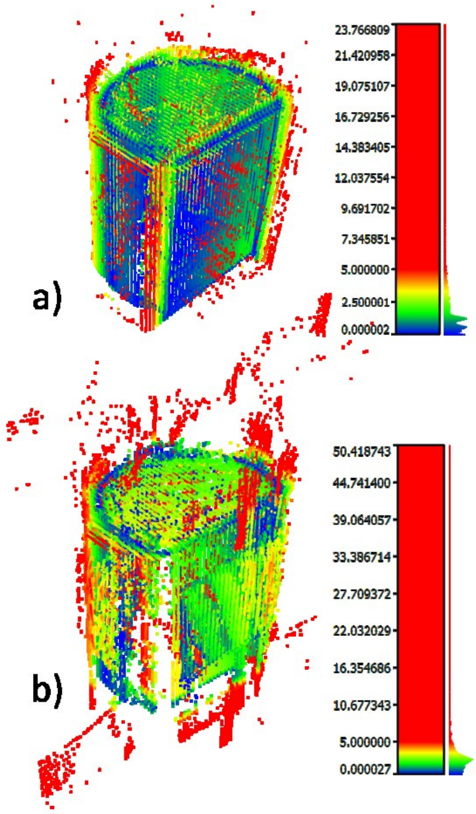

Fig. 14: SAFT (a) and TFM (b) derived surface point cloud reconstructions of the cylindrical sample.

The effect of the curved surfaces is detailed in Fig. 15, where, around the centre of the curve, there appears to be a complete lack of curvature. This is similar to the features seen in Fig. 11 (b), the cause of which was the specular effects explained in section 3.2. In addition to this, when the probe moved farther from the centre of the curve, it can be seen that the lines became erratic and sparse. This was due to the distorting effect described in section 3.2, which did not provide obvious surface profiles to be extracted. The effects of curvature were prominent due to the positioning of the probe relative to the surface. If a path was employed which maintained normality to the sample surface, the effects would be reduced significantly.

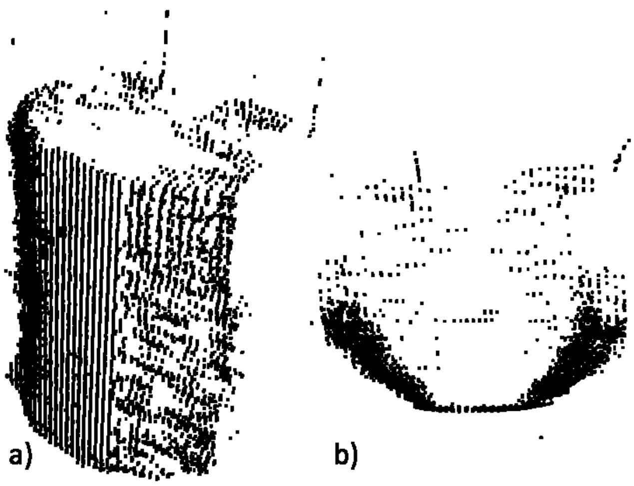

Fig. 15: Surface reconstruction of the scan line that moves past the side opposing the flat surface of the cylindrical sample.

Table 1: Results of comparison with the reference models.

Sample	Processing Method	Mean Error (mm)	Maximum Error (mm)	Standard Deviation (mm)
Cuboid	TFM	3.7	44.0	3.1
Cuboid	SAFT	1.3	32.0	1.8
Sphere	TFM	3.0	25.2	2.6
Sphere	SAFT	2.9	42.9	4.4
Cylinder	TFM	3.1	52.2	4.5
Cylinder	SAFT	2.0	26.5	2.1

3.4 Discussion

The images constructed using TFM have had noticeably higher side lobe content, leading to an increase in the associated measured surface error for all three sample geometries investigated. In traditional ultrasound imaging, the term "side lobes" refers to imaging artefacts that result from regions of ultrasonic energy which are produced off-axis relative to the main lobe during transmission [37]. In synthetic aperture methods, however, side lobes are regions of high intensity not attached to the main lobe [19] and are a result of the image construction algorithm itself.

This can be explained by considering a point spread function generated using TFM. Employing a ray-based model based on that described in [8], FMC data was simulated, providing the response of a point reflector 10 mm from the face of the probe. The medium was defined with a longitudinal speed of sound of 1480 ms^{-1}, the array with 32 elements with a pitch of 0.7 mm and the sampling frequency as 100 MHz. The output of each element was modelled as a 5 cycle, Gaussian windowed tone burst with a centre frequency of 5 MHz and a -6 dB bandwidth of 50%, as has been typically employed before [8], [38], [39].

As can be seen in Fig. 16 (a), as well as the obvious point reflector representation at (a,b), there are significant side lobes located either side. The reason these occur is because the algorithm does not discriminate based on pixel location. For example, for transmitting element i and receiving element j, the time of flight to pixel locations (c,d) – the position of the point reflector – and (a, b) is the same. In other words, with reference to (1), $S_{i,j}\left(T_{i(a,b)} + T_{j(a,b)}\right) = S_{i,j}\left(T_{i(c,d)} + T_{j(c,d)}\right)$. As such, the contribution to the intensity of the pixel, for that particular transmit-receive pair will be the same at both locations.

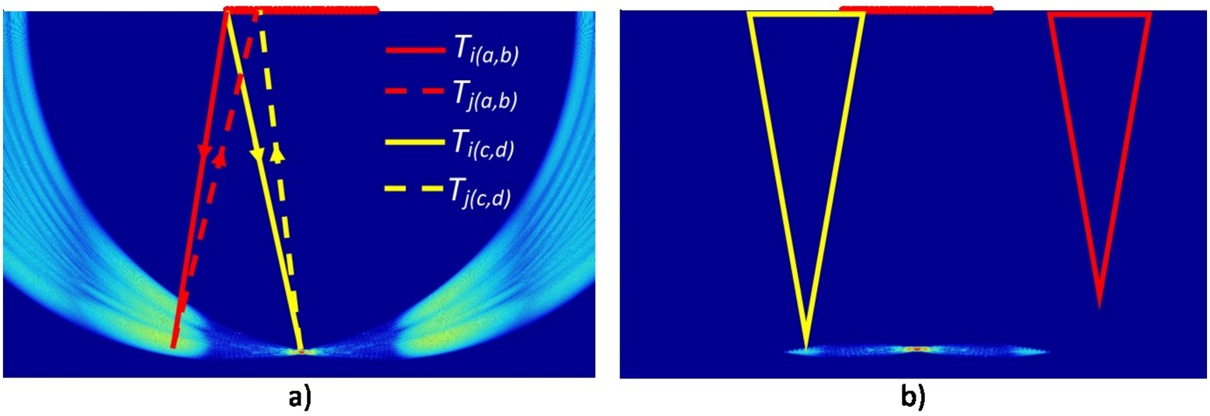

Fig. 16: Part (a) shows a TFM reconstruction of a simulated FMC of a point reflector. Element positions are displayed as circles, while the paths of travel for two pixels at (a, b) and (c, d) are shown for the same transmission and reception. Part (b) displays a SAFT reconstruction of the same data, showing the reconstructing triangles for two pixels.

This principle is repeated over the entire region shown as non-zero in Fig. 16 (a). Because reflections are not represented in A-scans as infinitely thin peaks, but instead have a width, the pixels which receive a non-zero contribution need not possess the exact time of travel as the position of the true reflector. Instead, the time of travel must only be within a range proximal to the true reflector. This is illustrated in Fig. 17, which shows the exact time of travel of the point reflector, P, along with the region in which the contribution will be non-zero, R. This leads to numerous non-zero contributions throughout the image, confined by the geometrical combinations of pixel position, transmit-receive pairings and reflector position. There are particular regions where more of the transmit-receive pairs and pixel positions meet the criteria for non-zero contribution. This leads to regions where the side lobes have a particularly high intensity, as can be seen in Fig. 16 (a).

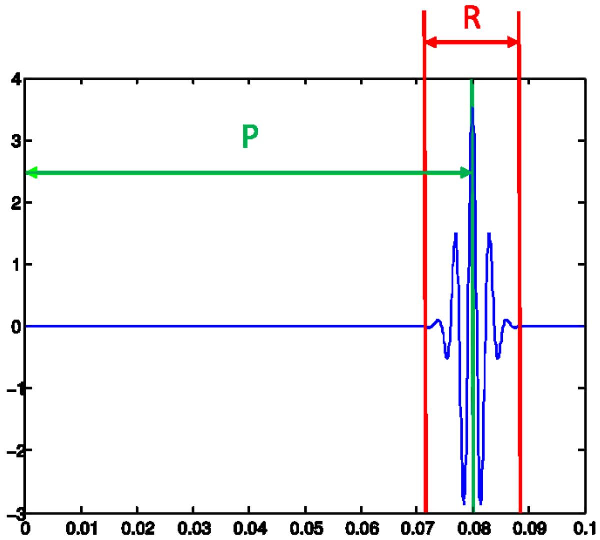

Fig. 17: An A-scan from the point reflector FMC, showing the time of travel for the point reflector, P, and the range in which pixel contributions will be non-zero, R.

The 20° SAFT reconstruction of the same FMC data is presented in Fig. 16 (b), displaying a significantly lower side lobe contribution. Also shown are the reconstructing triangles for two pixels, with the base determining the receiving elements employed reconstruction sub aperture. The left pixel, showing a yellow triangle, has included a number of elements in the reconstructing sub aperture, while the right pixel has none. The result of this is that the left pixel has a non-zero intensity, while the right has zero intensity.

While this example employed a point reflector for simplicity, the principle presented is valid for any physical reflector which elicits a high intensity response in a number of the receiving elements and, therefore, produces a region of high intensity in the reconstructed image. As such, the flat surface shown in Fig. 7 can be thought of as a densely populated line of discrete, strong reflections, each creating side lobes. These have merged to form a thick line above the true surface. This effect is significantly lessened in the SAFT images, as the number of possible transmit-receive pairs is limited by the fact that the number of elements considered in reception is significantly less than that in TFM. This geometrically restricts the regions in which side lobes can be formed, as demonstrated in Fig. 16 (b). Given that a narrow sub aperture has been employed in this study, the likelihood of side lobes is low. If, however, the sub aperture size was increased, the possibility of side lobes would increase.

It is pertinent also in this discussion to address the issue of path geometry used in this study and the likelihood of such path scanning to be employed in a real application on a knee joint. In the present study, the fact that the basic geometry of the test samples was known a-priori allowed an immediate construction of a suitable probe scan path. In the final application, the geometry is unknown and such simple path construction is considered unlikely to produce an outcome on a single pass. It is likely that a 2 (or more) stage scan would be required in practice, with an initial coarse scan used to generate a basic representation of the bone surface. This coarse scan would then be used to construct an optimised scan path around the knee joint for the subsequent high resolution scan.

4. Summary and Conclusions

The performance of two synthetic aperture methods to accurately image a number of surfaces corresponding to three precision metallic objects with surfaces including curves, flat surfaces and edges has been presented. The Total Focussing Method (TFM) and the Synthetic Aperture Focussing Method (SAFT) were selected for image reconstruction, as these represented extremes of the imaging approach, employing the full synthetic aperture width and a minimal aperture width, respectively.

The metallic samples of known geometries (cuboid, sphere and cylinder) were imaged using a 5 MHz, 128 element 1D phased array, which was manipulated around the samples using a precision robotic positioning system, allowing for repeatable ultrasound derived 3D surface point clouds to be created. A global thresholding technique was presented that allowed extraction of the surface profiles and these were compared with the known geometry samples to provide a quantitative measure of error of 3D surface reconstruction.

Producing mean errors of 1.3 mm, 2.9 mm and 2.2 mm using SAFT and 3.7 mm, 3.0 mm and 3.1 mm using TFM for the cuboidal, spherical and cylindrical samples respectively, SAFT offered significant improvements in accuracy over TFM. This was a result of improved clarity of surface representations, which allowed for more accurate surface profile extraction. The reduction in the width of the synthetic aperture of SAFT allowed for this, eliminating the side lobes associated with TFM. While the use of SAFT imposed a slight reduction in resolution and coverage, it provided mean errors approaching the resolution of CT - the 'gold standard' in preoperative imaging for robotic knee arthroplasty. Therefore, for unknown surface types, a narrow aperture SAFT is the superior imaging method, indicating that it would provide the most accurate depictions of the complex surfaces in the prescribed biomedical application. Additionally, this result has significant implications for dual-media time of flight correction techniques within NDT, in that the employment of a narrow aperture SAFT could allow for automatic identification of the surfaces of parts with more complex shapes than would be possible with TFM.

The results presented in this paper indicate that synthetic aperture methods are capable of highly accurate surface imaging, provided a narrow synthetic aperture is employed. While edges and curved surfaces were responsible for errors, the shapes employed were rudimentary and intended to amplify the associated effects. As such, they serve as a worst case scenario and the effects would be expected to be significantly lessened in bone surface imaging due to the increase in shape complexity and decrease in surface specularity.

In addition to the change in the surface type, the final application would also include multiple soft tissue interfaces preceding the bony surface under inspection. However, it is predicted this would not cause serious ill-effect in the resulting images, in that the acoustic impedance mismatch between various soft tissues and water is small relative to that between soft tissue and bone. This, however, can only be confirmed by experiment. Therefore, future investigations should concentrate on real bone surfaces with preceding soft tissue layers with an aim to achieve a closer representation of the bony surfaces found in the knee joint.

Acknowledgements

This work has been conducted with financial support from the EPSRC.

References

[1] M. Roche, P. F. O'Loughlin, D. Kendoff, V. Musahl, and A. D. Pearle, "Robotic arm-assisted unicompartmental knee arthroplasty: preoperative planning and surgical technique.," *Am. J. Orthop. (Belle Mead. NJ).*, vol. 38, no. 2 Suppl, pp. 10–5, Feb. 2009.

[2] W. Siebert, S. Mai, R. Kober, and P. F. Heeckt, "Technique and first clinical results of robot-assisted total knee replacement.," *Knee*, vol. 9, no. 3, pp. 173–80, Sep. 2002.

[3] J. Cobb, J. Henckel, P. Gomes, S. Harris, M. Jakopec, F. Rodriguez, a Barrett, and B. Davies, "Hands-on robotic unicompartmental knee replacement: a prospective, randomised controlled study of the acrobot system.," *J. Bone Joint Surg. Br.*, vol. 88, no. 2, pp. 188–97, Feb. 2006.

[4] T. Hananouchi, N. Nakamura, A. Kakimoto, H. Yohsikawa, and N. Sugano, "CT-based planning of a single-radius femoral component in total knee arthroplasty using the ROBODOC system," *Comput. Aided Surg.*, vol. 13, no. 1, pp. 23–29, 2008.

[5] H. L. Fred, "Drawbacks and Limitations of Computed Tomography," *Texas Hear. Inst. J.*, vol. 31, no. 4, pp. 345–348, 2004.

[6] J. Henckel, R. Richards, K. Lozhkin, S. Harris, F. M. Rodriguez y Baena, a R. W. Barrett, and J. P. Cobb, "Very low-dose computed tomography for planning and outcome measurement in knee replacement. The imperial knee protocol.," *J. Bone Joint Surg. Br.*, vol. 88, no. 11, pp. 1513–8, Nov. 2006.

[7] J. M. Albert, "Radiation risk from CT: Implications for cancer screening," *Am. J. Roentgenol.*, vol. 201, no. July, pp. 81–87, 2013.

[8] C. Holmes, B. W. Drinkwater, and P. D. Wilcox, "Post-processing of the full matrix of ultrasonic transmit–receive array data for non-destructive evaluation," *NDT E Int.*, vol. 38, no. 8, pp. 701–711, Dec. 2005.

[9] J. A. Jensen, S. I. Nikolov, K. L. Gammelmark, and M. H. Pedersen, "Synthetic aperture ultrasound imaging.," *Ultrasonics*, vol. 44 Suppl 1, pp. e5–15, Dec. 2006.

[10] M. Njiki, A. Elouardi, S. Bouaziz, O. Casula, and O. Roy, "A real-time implementation of the Total Focusing Method for rapid and precise diagnostic in non destructive evaluation," in *2013 IEEE 24th International Conference on Application-Specific Systems, Architectures and Processors*, 2013, pp. 245–248.

[11] Jie Zhang, B. W. Drinkwater, and P. D. Wilcox, "Comparison of ultrasonic array imaging algorithms for nondestructive evaluation.," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 60, no. 8, pp. 1732–45, Aug. 2013.

[12] A. J. Hunter, B. W. Drinkwater, and P. D. Wilcox, "The wavenumber algorithm for full-matrix imaging using an ultrasonic array.," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 55, no. 11, pp. 2450–62, Nov. 2008.

[13] P. D. Wilcox, C. Holmes, and B. W. Drinkwater, “Advanced reflector characterization with ultrasonic phased arrays in NDE applications.,” *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 54, no. 8, pp. 1541–50, Aug. 2007.

[14] A. Mcgilp, J. Dziewierz, T. Lardner, A. Gachagan, and C. Bird, “Inspection of Complex Components using 2D Arrays and TFM,” in *BINDT*, 2014.

[15] M. Weston, P. Mudge, C. Davis, and A. Peyton, “Time efficient auto-focussing algorithms for ultrasonic inspection of dual-layered media using Full Matrix Capture,” *NDT E Int.*, vol. 47, pp. 43–50, 2012.

[16] X. Han, W. Wu, P. Li, and J. Lin, “Application of Ultrasonic Phased Array Total Focusing Method in Weld Inspction Using an inclined Wedge,” in *Symposium on Piezoelectricity, Acoustic Waves, and Device Applications*, 2014, pp. 114–117.

[17] J. Camacho and J. F. Cruza, “Auto-Focused Virtual Source Imaging With Arbitrarily Shaped Interfaces,” *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 62, no. 11, pp. 1944–1956, 2015.

[18] P. Foroughi, E. Boctor, M. J. Swartz, R. H. Taylor, and G. Fichtinger, “Ultrasound Bone Segmentation Using Dynamic Programming,” *2007 IEEE Ultrason. Symp. Proc.*, pp. 2523–2526, Oct. 2007.

[19] J. Dziewierz, A. Gachagan, N. Lord, and A. J. Mullholland, “An application-specific design approach for 2D ultrasonic arrays .,” in *51st Annual Conference of the British Institute of Non-Destructive Testing 2012 (NDT 2012)*, 2012.

[20] J. Lambert, P. Antoine, G. Gens, F. Bimbard, L. Lacassagne, and E. Iakovleva, “Performance Evaluation of Total Focusing Method on GPP and GPU,” in *2012 Conference on Design and Architectures for Signal and Image Processing (DASIP)*, 2012, pp. 1–8.

[21] M. A. Conditt and M. W. Roche, “Minimally Invasive Robotic-Arm-Guided Unicompertmental Knee Arthroplasty,” *J. Bone Jt. Surg.*, vol. 91, no. 1, pp. 63–68, 2009.

[22] M. Citak, E. M. Suero, M. Citak, N. J. Dunbar, S. H. Branch, M. a. Conditt, S. a. Banks, and A. D. Pearle, “Unicompartmental knee arthroplasty: Is robotic technology more accurate than conventional technique?,” *Knee*, vol. 20, no. 4, pp. 268–271, 2013.

[23] B. Davies, M. Jakopec, S. J. Harris, F. Rodriguez y Baena, A. Barrett, A. Evangelidis, P. Gomes, J. Henckel, and J. Cobb, “Active-Constraint Robotics for Surgery,” *Proc. IEEE*, vol. 94, no. 9, pp. 1696–1704, 2006.

[24] D. I. A. Lines, I. G. Pettigrew, K. J. Kirk, S. Cochran, and J. A. Skamstad, “Rapid Distributed data Collection and Processing With Arrays - the Next Step Beyond Full Waveform Capture,” *9th Jt. FAA/DoD/NASA Aging Aircr. Conf.*, pp. 1–11, 2006.

[25] R. H. Brown, S. G. Pierce, I. Collison, B. Dutton, J. Jackson, T. Lardner, C. Macleod, and M. Morozov, “Automated Full Matrix Capture for Industrial Processes,” *41st Annu. Rev. Prog. Quant. Nondestruct. Eval.*, vol. 34, pp. 1967–1976, 2015.

[26] R. Halmshaw, *Non-Destructive Testing*, 2nd ed. London: Edward Arnold, 1991.

[27] A. Maurer, W. Deodorico, R. Huber, and T. Laffont, “Aerospace Composite Testing Solutions using Industrial Robots,” *18th World Conf. Nondestruct. Test.*, no. April, p. 7, 2012.

[28] F. Mohr, W. Haase, and D. Ungerer, “Automated Non-Destructive Examination of Complex Shapes,” in *11th European Conference on Non-Destructive Testing*, 2014, no. ECNDT.

[29] P. Louviot, A. Tachattahte, and D. Gardener, “Robotised UT Transmission NDT of Composite Complex Shaped Parts,” *NDT Aerosp. 2012 - We.3.B.2*, vol. 4th Intern, pp. 1–8, 2012.

[30] E. Cuevas and S. Hernández, “Robot-based solution To Obtain an Automated , Integrated And Industrial Non-Destructive Inspection Process,” in *6th International Symposium on NDT in Aerospace*, 2014, no. November, pp. 12–14.

[31] C. Mineo, S. G. Pierce, P. I. Nicholson, and I. Cooper, “Robotic path planning for non-destructive testing – A custom MATLAB toolbox approach,” *Robot. Comput. Integr. Manuf.*, vol. 37, pp. 1–12, 2016.

[32] KUKA, *KUKA System Software 5*. 2010.

[33] KUKA, *KUKA Robot Sensor Interface 2.3*. 2009.

[34] R. Summan, S. G. Pierce, C. N. Macleod, G. Dobie, T. Gears, W. Lester, P. Pritchett, and P. Smyth, “Spatial calibration of large volume photogrammetry based metrology systems,” *Measurement*, vol. 68, pp. 189–200, 2015.

[35] D. Romero, O. Mart, and C. J. Mart, “Using GPUs for beamforming acceleration on SAFT imaging,” in *2009 IEEE International Ultrasonics Symposium Proceedings*, 2009, pp. 1334–1337.

[36] P. D. Wilcox, C. Holmes, and B. W. Drinkwater, “Enhanced Defect Detection and Characterisation by Signal Processing of Ultrasonic Array Data,” *ECNDT 2006*, pp. 1–9, 2006.

[37] P. Y. Barthez, R. Léveillé, and P. V Scrivani, “Side lobes and grating lobes artifacts in ultrasound imaging.,” *Vet. Radiol. Ultrasound*, vol. 38, no. 5, pp. 387–93, 1996.

[38] M. Li and G. Hayward, “Ultrasound nondestructive evaluation (NDE) imaging with transducer arrays and adaptive processing.,” *Sensors (Basel).*, vol. 12, no. 1, pp. 42–54, Jan. 2012.

[39] C. Fan, M. Caleap, M. Pan, and B. W. Drinkwater, “A comparison between ultrasonic array beamforming and super resolution imaging algorithms for non-destructive evaluation.,” *Ultrasonics*, vol. 54, no. 7, pp. 1842–1850, Jan. 2014.