
Real-time Freehand 3D Ultrasound Calibration

P-W. Hsu, R. W. Prager
A. H. Gee and G. M. Treece

CUED/F-INFENG/TR 565

September 2006

University of Cambridge
Department of Engineering
Trumpington Street
Cambridge CB2 1PZ
United Kingdom

Email: pwh24@cam.ac.uk, rwp/ahg/gmt11@eng.cam.ac.uk

Real-time Freehand 3D Ultrasound Calibration

Po-Wei Hsu, Richard W. Prager, Andrew H. Gee and Graham M. Treece

University of Cambridge
Department of Engineering
Trumpington Street
Cambridge CB2 1PZ

Abstract

Z-fiducial phantoms allow 3D ultrasound probe calibration with a single B-scan. One of the main hindrances of using this phantom is the necessity of segmenting the wires reliably, which requires human intervention. In this paper, we have shown how we can solve this problem by mounting a thin rubber membrane on top of the phantom. The membrane is segmented automatically and the wires can be easily located as they are at known positions relative to the membrane. This enables us to segment the wires automatically at the full PAL frame rate of 25Hz, to produce calibrations in real-time, while achieving accuracies similar to those reported in the literature. We also devise a technique to improve the estimation of the elevational offset by capturing a few images of the planar membrane. If spatial calibration is known, fully automatic wire segmentation allows the fiducials to be tracked in real-time. This also enables temporal calibration to be performed in real-time as the probe is moved away from the phantom.

1 Introduction

Freehand three-dimensional (3D) ultrasound (Fenster et al., 2001) is a medical imaging modality with many applications in surgical navigation, radiotherapy planning and complex anatomy visualization and analysis (Gee et al., 2003). As the ultrasound probe is swept over the anatomy, the attached position sensor records the trajectory of the probe. The 3D volume of the anatomy can be constructed by matching the ultrasonic data with its corresponding position in space. There are however two main prerequisites that have to be overcome when using this imaging technique. They are (1) probe (spatial) calibration and (2) temporal calibration (Mercier et al., 2005).

1.1 Spatial Calibration

The position sensor that is attached to the ultrasound probe will measure the 3D location of the sensor, rather than the scan plane, relative to an external world coordinate system as shown in Figure 1. It is therefore necessary to find the position and orientation of the scan plane with respect to the electrical centre of the position sensor. This rigid-body transformation is determined through a process called probe (spatial) calibration.

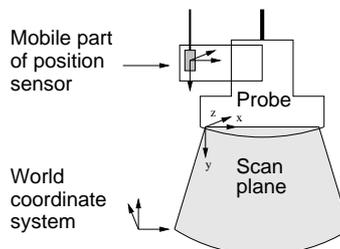


Figure 1: The coordinates associated with the scan plane and mobile part of the position sensor.

Performing spatial calibration by scanning an object with known geometric properties (a phantom) has been a research topic for many years. The simplest phantom is a point target formed by

a pair of cross-wires (Detmer et al., 1994; Gooding et al., 2005) or a small spherical ball bearing (Amin et al., 2001). In both cases, the point is scanned from different positions and its location marked in each B-scan. Each segmented point can be mapped to the sensor’s coordinate system using an assumed calibration, and then to the world coordinate system using the position sensor readings. If the calibration is correct, every point should be mapped to the same location in space. This observation can be expressed as an over-determined system of constraints on the calibration parameters that can be solved using optimisation techniques.

The main drawback of using this simple phantom is the amount of time required to perform one calibration. Automatic segmentation of isolated points in ultrasound images is seldom reliable. This is evident from the fact that many authors manually segment their targets in the B-scans (Detmer et al., 1994; Amin et al., 2001; Dandekar et al., 2005). This makes calibration a time consuming process. Secondly, it is very difficult to align the scan plane with the point target. The finite thickness of the ultrasound beam causes the point target to appear in the B-scan even if the target is not in the elevational centre of the scan plane. This error can be up to a few millimetres in the elevational direction, depending on the ultrasound probe and the skill of the user. Furthermore, it is important that the point target has been scanned in a non-degenerate configuration, so that the resulting system of constraints will not be under-determined such that no unique solution exists (Prager et al., 1998).

As freehand 3D ultrasound gained popularity over the last decade, a lot of research has been undertaken to make the calibration process more rapid and reliable, sometimes at the expense of using a more complicated phantom. Prager et al. (1998) was the first to scan a plane instead of the point target. The design complexity of the plane ranges from the floor of a water tank (Prager et al., 1998), a plexiglass plate (Rousseau et al., 2005), a nylon membrane (Langø, 2000) to a precision made Cambridge phantom (Prager et al., 1998). In each case, the plane will appear as a straight line in the B-scans. This can be reliably segmented automatically to allow a rapid calibration. The plane still needs to be scanned in a deterministic configuration, but the eigenvalue metric may be used to assess whether the plane has been scanned from a sufficiently diverse range of positions (Hsu et al., 2006).

A 3D localizer (pointer), which can be used to locate fixed points in space, is often used to aid spatial calibration. Muratore and Galloway Jr. (2001) performed spatial calibration by locating points directly in the scan plane by using a pointer. Since the located points are known in the two coordinate systems, the corresponding transformation can be solved using various algorithms (Bouchet et al., 2001). For the same reason when using a point target phantom, it is difficult to align the scan plane with the tip of the pointer, which is essentially a point target. Khamene and Sauer (2005) improved on this technique by imaging a rod transversely. Both ends of the rod acts as 3D localizers, thereby defining the orientation of the rod in space. Again, each image of the rod sets up a constraint on the calibration parameters. Spatial calibration is then found by using optimization techniques. These techniques however require two targets to be tracked simultaneously during spatial calibration.

A different class of phantoms is the two-dimensional (2D) alignment phantom. These phantoms have the advantage that only one frame is needed for spatial calibration. Sato et al. (1998) aligned the scan plane with a thin board. Three corners of the board are located in space using the aforementioned pointer. These corners are segmented manually in the ultrasound B-scan, thereby solving for the spatial calibration parameters. There are two main difficulties when using this phantom. Firstly, it is very difficult to align the phantom precisely with the scan plane. Secondly, the manual segmentation of the corners means that calibration is a lengthy process.

Gee et al. (2005) designed a mechanical device where the 2D phantom is aligned using micrometers. They mounted wedges onto the planar phantom so that any misalignment could be easily detected in the B-scans. These wedges also served as the fiducial points to be located in space, which are semi-automatically segmented in the B-scans. The phantom itself is not located in space using a pointer. Instead a gantry is designed so that both the position sensor and the probe can be precisely mounted onto the phantom.

The Z-fiducial phantom was designed to solve the alignment problem in 2D phantoms (Comeau et al., 1998). The intersection of the Z-fiducials with the scan plane forms a virtual 2D phantom

that can be used to define the position of the scan plane in space. The phantom itself is defined in space either by using a pointer (Pagoulatos et al., 2001) or using another position sensor (Bouchet et al., 2001; Lindseth et al., 2003). The main drawback of using this type of phantom is the difficulty associated with segmenting isolated points in each B-scan. Furthermore, each wire does not appear as a single dot or disc in the B-scans, but rather as a smeared ellipse, compounding the segmentation difficulty. Nevertheless, when given a predefined search region, automatic segmentation is possible (Lindseth et al., 2003).

1.2 Temporal Calibration

In modern freehand 3D ultrasound systems a computer is usually used to synchronize data from the ultrasound machine and the position sensor (Prager et al., 1999; Meairs et al., 2000; Huang et al., 2005). The computer sends requests to the associated hardware and time-stamps the ultrasonic and position data when it arrives, as shown in Figure 2. However, these data arrive with an unknown delay after having been generated by the hardware. Temporal calibration involves finding the relative delay between the two data streams so that matching B-scans and positions are available for further analysis.

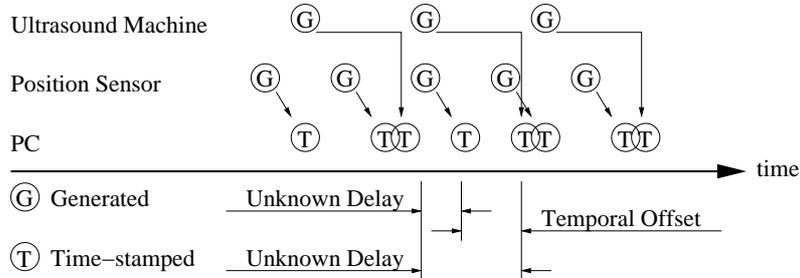


Figure 2: Both the ultrasound machine and the position sensor generate (G) data at their own frequencies. They are received and time-stamped (T) by the PC after an unknown delay. Temporal calibration finds the relative delay between the two data streams.

During both spatial and temporal calibration, the data from the ultrasonic and position streams are monitored and analyzed while scanning the selected phantom. It is therefore necessary to decide beforehand which calibration is to be performed first. If the temporal offset is to be computed first, the exact location of the phantom in space is unknown during temporal calibration (Prager et al., 1999; Huang et al., 2005). If, on the other hand, the spatial calibration is to be performed first, continuous scanning of the phantom will be subject to temporal distortions. A common practise to by-pass this problem is to capture images of the phantom while holding the probe stationary for a short duration at each position (Gooding et al., 2005; Hsu et al., 2006).

The simplest way of finding the temporal offset between the ultrasonic and position data streams, without utilizing any spatial information, is to apply and detect a step input in the two data streams (Prager et al., 1999; Meairs et al., 2000). This may be achieved by holding the probe against the skin and then suddenly jerking the probe away from the body. The required offset is the difference between the two time-stamps in the two streams where the abrupt change took place. The accuracy of this technique is however only $\pm \frac{t+T}{2}$, where t and T are the temporal resolutions of the two data streams. Nevertheless, if the ultrasonic frames are captured at the full PAL frame rate at 25Hz, and the position data exceeds this rate, this simple temporal calibration technique is sufficient for clinical volumetric measurements (Rousseau et al., 2006).

A more accurate technique is to scan a fixed target, such as a wire (Burcher, 2002) or a plane (Treece et al., 2003), while the probe is moved in one direction. The target in each B-scan is segmented. Since the probe is only moved in one direction, the location of the target in space can be approximated from the ultrasound images without spatial calibration. Temporal calibration is found where there is maximum correlation between the two data streams. Instead of approximating

the location of the target plane in space, Rousseau et al. (2006) parameterized the plane images in each B-scan and estimated the temporal misalignment by matching the position sensor data with the line parameters.

In contrast to the aforementioned techniques where temporal calibration is performed before spatial calibration, spatial calibration may be performed in certain circumstances without knowing the temporal offset. Once the spatial calibration is known, then any target that has been scanned can be located in space. Nakamoto et al. (2003) was the first to follow this approach by scanning a point target. Gooding et al. (2005) applied image alignment techniques and scanned a phantom with two distinct materials. The correct temporal offset occurs when the reconstructed image of the phantom best resembles its known geometric properties.

As we have seen, many research groups have developed new techniques to make spatial calibration faster and easier to perform. However, without fully automatic segmentation algorithms and a suitable phantom that allows calibration with only one or a very small number of frames, calibration typically takes a few minutes to complete. In this paper, we have designed a novel phantom based on the Z-phantom (Bouchet et al., 2001; Pagoulatos et al., 2001; Lindseth et al., 2003). Our innovation was inspired by Langø (2000), who used a thin membrane as a plane based phantom. We realized that the membrane does not just produce a strong reflection, but the ultrasound waves are barely distorted by the membrane. As a result, the structures beneath the membrane appear clearly in the B-scans. In this paper, we show how we have used the membrane as a landmark to enable fully automatic segmentation in a Z-phantom, thus producing calibrations in real-time. This means that we can monitor changes in the volume of a prerecorded anatomy as calibration is performed. Re-calibration due to a change in depth or zoom will only take a few seconds, provided that the phantom remained stationary, allowing the clinician to change the settings as desired during medical diagnosis.

2 The Calibration Phantom

Figure 3(a) shows the calibration phantom. The phantom consists of two parallel polyacetal blocks 5cm apart. We have drilled a number of holes with 0.5mm diameter in these blocks in a predefined configuration. These holes serve as the locations for the ‘Z’ shaped fiducials, which are placed in 5 rows with 4 fiducials in each row, except the first row which has 3 fiducials, as shown in Figure 3(b). A 1mm thick translucent 40° Shore A silicone rubber membrane is clamped under tension on top of the phantom. There are 3 cone-shaped divots on the two blocks for fitting our 3mm ball-pointed 3D localizer. These divots serve as the principal axes of the phantom and are used to define the phantom in 3D space. All dimensions are precision manufactured by our workshop with a tolerance of $\pm 0.1\text{mm}$. We then interwove an ordinary off-the-shelf fishing line with a diameter of approximately 0.5mm through the holes in the two blocks to form the ‘Z’ shaped fiducials.

In order to perform calibrations in real-time, we need to minimize the computation required during (1) segmentation and (2) solving for the calibration parameters.

The time needed for segmentation is reduced by the introduction of the membrane, which appears as a straight line in the B-scans, and can be segmented automatically (Prager et al., 1998). The positions of the wires are at known depths beneath the membrane by construction. If the image scales are known, the search region for the wires can be limited to five narrow strips, one at each depth, as shown in Figure 4. Since the fiducials are mounted parallel to the rubber membrane, the algorithm remains robust even if the probe is at an inclination to the phantom. We have used a fairly simple algorithm to segment the smeared blobs. Since the smeared image of the wire appears to be approximately 1–3mm in diameter, the average intensities of each $1\text{mm} \times 1\text{mm}$ block in the search region are computed. A user defined threshold is placed on the average intensities. This divides the search region into a number of connected regions. Potentially, each region represents a smeared image of the wire. The maximum average intensity of each connected region is located. The centroid of a small rectangular region (approximately $0.5\text{mm} \times 0.25\text{mm}$) centred at this location is computed and stored as a potential wire location. Since we have limited

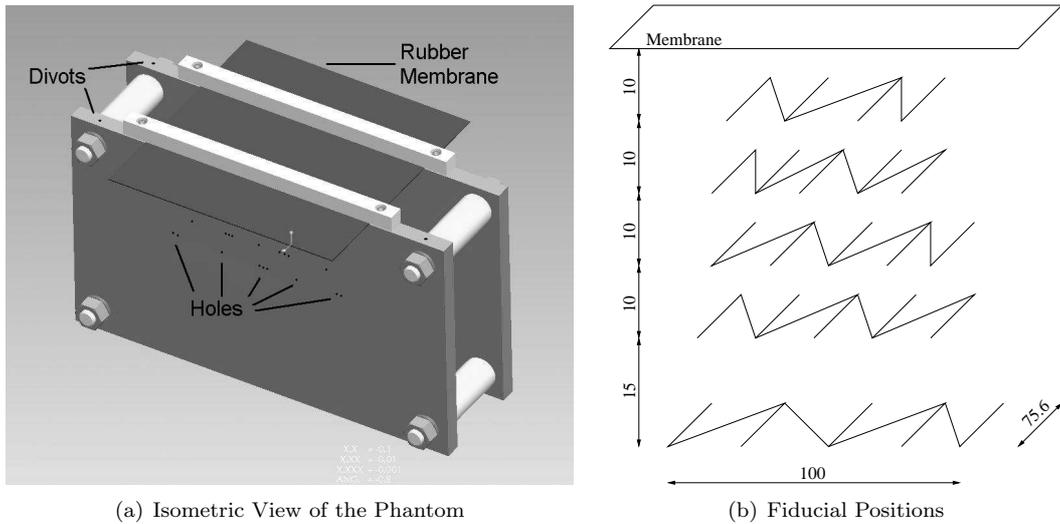


Figure 3: The calibration phantom.

the search region to a narrow strip, our simple segmentation algorithm is robust over a wide range of thresholds, with only a few falsely detected potential wire locations. These falsely detected wires are easily rejected by fitting the potential wire locations to the construction geometry of the wire configuration.

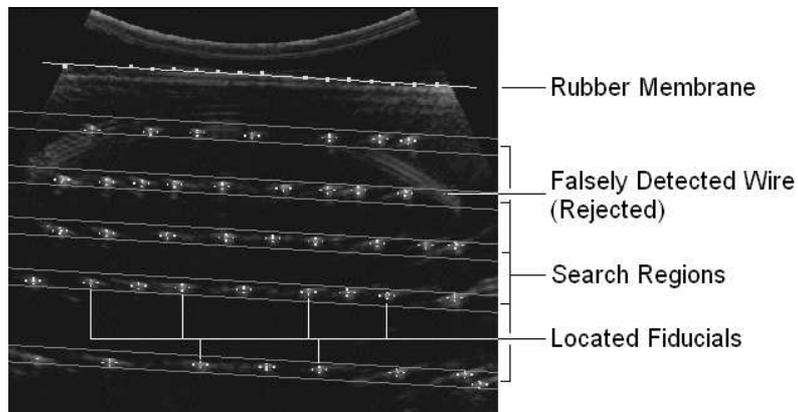


Figure 4: Wire segmentation of a typical B-scan image. The search regions are limited to five narrow strips at known depths beneath the rubber membrane.

In general, spatial calibration needs to find the horizontal and vertical scales in the B-scan images and the 6 parameters that define the phantom in space in addition to the 6 parameters that define the B-scan relative to the sensor's position. The values of the first 8 variables can be found using a separate protocol prior to phantom scanning. We first compute the scale values using the distance measurement tool provided by the ultrasound machine (Hsu et al., 2006). A calibrated pointer is then used to locate the phantom in space by its divots. This leaves only the 6 parameters that define the rigid body transformation from the scan plane to the position sensor to be found by calibration.

In our phantom design, the wires are interwoven through the holes H_1, H_2, H_3, H_4, H_5 and H_6 in a 'Z' shape, as shown in Figure 5. These holes are precision manufactured relative to the divots and are therefore known in space. The 3D positions of A, B, C and D can be calculated using trivial coordinate geometry. When the probe is placed over the wire configuration, the scan plane

intersects the wires at M , Z and N . These points can be located in the B-scan image. Since the image scales have been determined, the distances $|MZ|$ and $|MN|$ can be measured off the B-scan image. The location of Z can thus be computed: $\mathbf{Z} = \mathbf{B} + \frac{|BZ|}{|BC|}(\mathbf{C} - \mathbf{B}) = \mathbf{B} + \frac{|MZ|}{|MN|}(\mathbf{C} - \mathbf{B})$, since $\triangle BMZ$ and $\triangle CNZ$ are similar.

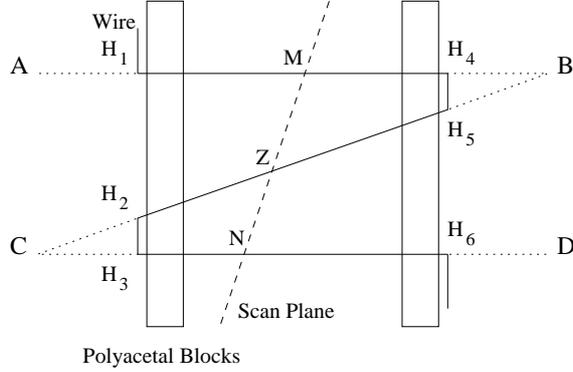


Figure 5: The top view of a Z-fiducial intersecting the scan plane.

For each ‘Z’ shape wire configuration, the Z-fiducial (labelled Z in Figure 5) is located in space using the aforementioned equation. This can be transferred to the sensor’s coordinate system using the inverse of the position sensor’s readings. We now post-process the previously segmented fiducial on the B-scan images to allow for the sound speed distortions caused by the rubber membrane (Achenbach, 1973). Since we have previously segmented each Z-fiducial in the B-scans, we have located the fiducial in the B-scan and the sensor’s coordinate systems. Assuming we have located at least 3 non-linear fiducials, we can compute the corresponding calibration that best fits the two data sets (Arun et al., 1987).

3 Spatial Calibration

The main advantage in the introduction of a membrane is to allow fully automatic segmentation and to output spatial calibrations in real-time. Segmentation is therefore highly dependent on the image quality of the B-scans. In order to demonstrate the robustness of our phantom to poor image qualities, we have chosen to calibrate our two-generation old Toshiba¹ model SSA-270A/HG 3.75MHz curvilinear probe at a depth of 8cm and 15cm. The probe has served many years in a clinical environment, followed by much abuse in our laboratory, have dead crystals and produces very poor images. One focus was used at the middle of the image in order to produce B-scans as fast as possible at 18Hz. The B-scan images were digitized using a Brooktree² BT878 frame-grabber card and transferred to a 3.0GHz PC running Linux. An area of 494×420 pixels of the B-scan images was cropped for calibration. The probe was tracked using an AdapTrax³ infrared LED target for the Polaris⁴ optical tracking system.

The calibration was performed by submerging our phantom in water at approximately 48 Celsius to match the speed of sound in water and soft tissue. The scales in the B-scan images were computed using the distance measurement tool available on the ultrasound machine (Hsu et al., 2006). The phantom was located in space by its divots using a pre-calibrated pointer. The probe was held perpendicularly over the phantom about 1–2cm above the rubber membrane. The B-scan produced is shown in Figure 4. Despite the overall noisy image, the segmentation algorithm has segmented on average 14 of the 15 fiducials that appears on the B-scan images when scanning at

¹Toshiba Corporation, <http://www.toshiba.com/>

²Conexant Systems Inc., <http://www.conexant.com/>

³Traxtal Technologies, <http://www.traxtal.com/>

⁴Northern Digital Inc., <http://www.ndigital.com/>

8cm, and segmented 18 out of the total 19 fiducials when calibrating at 15cm. Any undetected fiducial point is simply excluded while solving for the calibration parameters.

3.1 Precision and Accuracy

We now need to examine the accuracy of the calibration that is produced by using this phantom. We therefore scan the phantom 100 times. The probe is removed and reinserted into the water-bath between each scan. During each repetition, one frame is saved with the corresponding sensor information, phantom location and the calibration results for accuracy assessment.

We divided our 100 frames into 10 groups with 10 frames each. One calibration was computed from each group, forming the 10 repetitions that we used to evaluate calibration consistency. Single frame calibrations were computed by taking the first frame from each group. Similarly, n frame calibrations were computed by finding the parameters that best fit all the fiducials segmented from the first n frames of each group.

Each point in the B-scan can be located in the sensor’s coordinate system by using the parameters found by calibration. The precision of calibrating with this phantom is assessed by examining the variation of specific points in the scan plane while located in the sensor’s coordinate system. The variation will be dependent on the points chosen in the scan plane. We chose the four corners and the centre of the B-scan as a representative set of points. For each of these five points, we compute their mean 3D location $\bar{\mathbf{x}}_i = \frac{1}{10} \sum_{j=1}^{10} \mathbf{x}_i^j$, where \mathbf{x}_i^j is the 3D location of point i mapped using the j -th calibration, with $1 \leq i \leq 5$ and $1 \leq j \leq 10$. The mean 3D error of each representative point is found by computing the mean distance of each representative point from their mean location ($\frac{1}{10} \sum_{j=1}^{10} |\mathbf{x}_i^j - \bar{\mathbf{x}}_i|$). The mean 2D error is also computed. This is the error in the scan plane alone, neglecting the error in the elevational direction.

Precision measures the reproducibility of the calibration. This does not directly measure the accuracy of the calibration. One technique for assessing the accuracy of the system is to scan an object, reconstruct its image in 3D space and compare distances or volume measurements from the reconstructed image with the known dimensions. Pagoulatos et al. (2001) and Lindseth et al. (2003) used their Z-phantoms to provide fiducials with known locations in 3D space. We therefore followed their approach and independently scanned our phantom and compared the 3D locations of the located fiducials computed by the previously recorded 10 calibrations and the locations computed from the phantom construction.

Figure 6 shows the accuracy and precision of our calibrations as more frames are included in each calibration. The error for the precision graph is the mean 3D error of the five representative points’ locations.

We compare our precision and accuracies based on 10 frame calibrations with Pagoulatos et al. (2001), where a similar number of frames and fiducials were used. We can deduce from Figure 6 that our accuracy saturates just below 1.5mm, slightly better than the 2.8mm error reported by Pagoulatos et al. (2001). The precision of our system has a mean 3D error of 0.8mm and 1.5mm when calibrating at 8cm and 15cm respectively. This error is slightly better than the errors reported by Pagoulatos et al. (2001), which ranged mostly between 1mm and 2mm depending on the location in the B-scan. Bouchet et al. (2001) and Lindseth et al. (2003), who attached position sensors to their phantoms and thereby eliminate errors caused by phantom movements during calibration, reported slightly higher accuracy by a fraction of a millimetre. This means that our system can produce calibrations at least comparable to those in the literature by including sufficient frames in each calibration. Since we are able to segment the fiducials automatically, a large number of frames can be acquired and processed in a few seconds.

From Figure 6 it can be seen that the error saturates if more than 5 frames are used to compute the calibration parameters. We will use the results from the 10 frame calibrations for error analysis. The variation of the five representative points in the B-scan as a result of calibration imprecision is shown in Table 1. The mean 2D and 3D error, as well as the 95% confidence interval is given in the same table.

From Table 1 it can be deduced that the 3D error is much larger than the 2D error. This means that the main component of the error is in the elevational direction. The reason is because

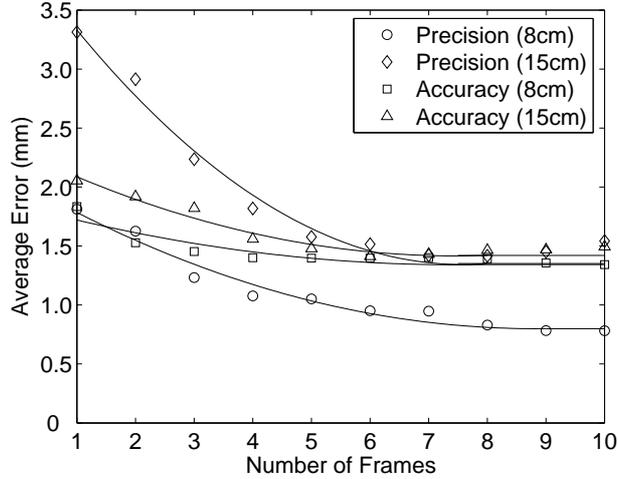


Figure 6: Variation of the error in terms of accuracy and precision as more frames are used in each calibration. Accuracy is calculated from the mean error of the located fiducials, and precision is calculated from the mean error of the 5 representative points’ locations.

Table 1: Variation in the four corners and centre of the B-scan due to calibration imprecision. The dimensions are in millimetres.

Variation in the four corners and centre of the B-scan due to calibration imprecision. The dimensions are in millimetres.

Depth	Position in the B-scan	Mean Error		95% Confidence	
		2D	3D	2D	3D
8cm	Top Left	0.41	0.74	< 0.76	< 1.29
	Top Right	0.37	0.82	< 0.69	< 1.43
	Bottom Left	0.43	0.94	< 0.79	< 1.64
	Bottom Right	0.37	0.95	< 0.69	< 1.66
	Centre	0.37	0.47	< 0.69	< 0.83
15cm	Top Left	0.54	1.74	< 1.00	< 3.04
	Top Right	0.53	1.09	< 0.98	< 1.90
	Bottom Left	0.54	1.51	< 1.00	< 2.63
	Bottom Right	0.55	2.31	< 1.02	< 4.06
	Centre	0.53	1.07	< 0.97	< 1.87

the fiducial’s position in the elevational direction is estimated from the distance between the wires in the B-scan images, and a slight error of a few pixels during wire segmentation will translate into an error of several millimetres in the elevational direction. Nevertheless, this error can be minimized by including multiple frames in each calibration.

3.2 Improving the Elevational Error

The relatively poor estimation of the calibration in the elevational direction is a characteristic weakness of the Z-fiducial phantom. This is evident, although not emphasized, from the calibrations performed in the literature with the Z-phantom (Pagoulatos et al., 2001). We therefore attempt to improve the elevational estimation.

We first capture a series of images of the planar membrane that is mounted on the phantom from

an oblique angle, as shown in Figure 7. It is important that the probe be tilted at approximately the same oblique angle on both sides when scanning the membrane, otherwise the beam-width effect would cause the membrane to appear at an incorrect depth (Prager et al., 1998). Since the membrane is mounted at a precise location on the phantom, the top of the membrane is known in space.

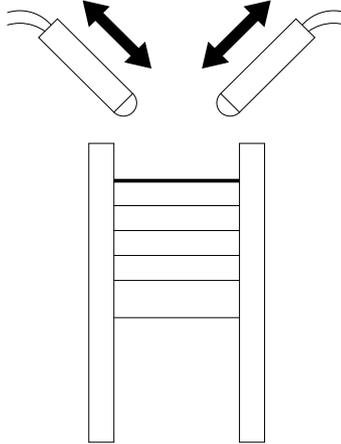


Figure 7: The motions used to scan the membrane as a plane to improve the elevational error.

As we scan the phantom as previously described, we obtain an preliminary calibration where the elevational offset is inaccurate. We will assume that the other 3 rotations and 2 translations are correct and independent of the elevational offset. This assumption is true if the elevational error is small. We form a new calibration by replacing the elevational offset with a variable, while holding the other parameters invariant. Now, points on the membrane can be mapped to the sensor’s coordinate system by the transformation with the unknown elevational offset, with each coordinate expressed as a linear function of the unknown offset. This point can be mapped to the world coordinate system by the sensor’s readings. Again, each of its coordinates would be a linear function of the unknown. This point must lie on the plane, whose equation is known. We have therefore obtained a linear equation with one unknown for each point on the plane. Every line image of the plane will give rise to two independent equations, hence n images of the plane will result in a system of $2n$ linear equations with one unknown, and the elevational offset is then solved by least square minimization.

We captured 16 images of the plane and replaced the elevational offset with the value calculated using the above method. Figure 8 shows that both 3D precision and accuracy are improved slightly, with the error saturating after 8 images. The relatively constant precision of the calibrations performed at 8cm depth suggests that the error has saturated and our correction did not worsen this precision.

4 Temporal Calibration

As we scan our phantom, the Z-fiducials can be located in space using two different methods. The first approach is to use the geometry of the phantom and the distances between the segmented wires as depicted in Figure 5. Alternatively, the fiducials can be located in the sensor’s coordinate system if the spatial calibration is known. Now, if there are no temporal misalignments between the position and ultrasonic data, we can locate the same fiducials in space by using the sensor’s readings. The position of the fiducials located in space using the aforementioned methods should be identical. If the probe is moved in some direction as we scan our phantom, temporal misalignment between the two data streams would be introduced. Furthermore, the error between the located fiducials with their true locations should relate positively to the amount of temporal distortion.

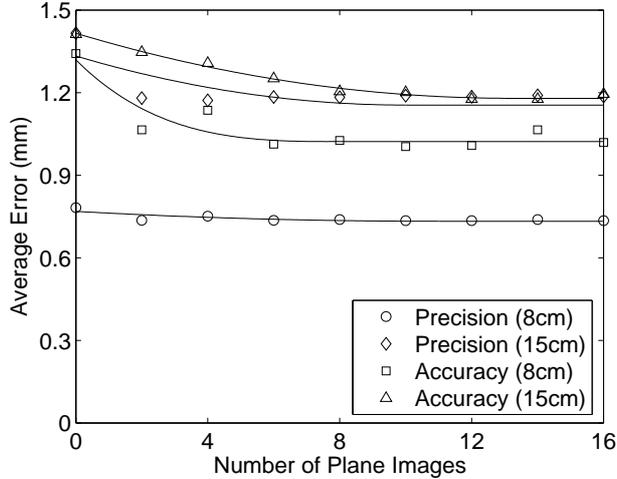


Figure 8: Precision and accuracy are improved due to a more accurate estimate of the elevational offset as more frames of the membrane are used.

We therefore locate the fiducials in 3D space using different position sensor readings and find the reading that best matches the fiducials with their true location in space. The accuracy of this approach will be dependent on spatial calibration, segmentation reliability and phantom construction tolerance.

With this in mind, we compute a spatial calibration by using all the fiducials obtained in the previous section. The probe is initially placed at about 1–2 cm above the rubber membrane, and then moved slowly by hand away from the phantom. As the probe is moved, the wires are segmented in real-time at 25 frames per second. At each instance, we are able to locate the segmented fiducials in space using the phantom geometry as described. At the same time, the same fiducials are located in the sensor’s coordinate system using the computed spatial calibration. The fiducials are then located in space using different position sensor readings that arrived at different instances in time from the position data buffer. If the temporal offset is correct, the located fiducial in space should match their true location determined by the phantom geometry.

Figure 9 shows a typical instance where the mean error of the fiducials located using the position sensor readings at different times relative to the time-stamp of the current frame. Since the probe is moved smoothly, we interpolate our data points with a cubic spline. The best temporal offset is assumed to be where the minimum of the interpolated curve occurs. A negative value signifies that the ultrasonic images are lagging the position data.

At each time instance during which the probe was moved, a curve similar to the one in Figure 9 can be plotted and a temporal offset calculated. We scanned our phantom 10 times at 15cm. Each time moving the probe for approximately 3cm away from the phantom. We take the temporal offset to be the mean of all the minima found at each time instance during each sweep. The temporal misalignment is found to be $-104 \pm 7.7\text{ms}$ (mean \pm standard deviation). The relatively low standard deviation of our temporal calibration compares favourably with the literature, where 5ms (Treece et al., 2003), 130ms (Gooding et al., 2005) and 18ms (Rousseau et al., 2006) were achieved.

5 Conclusion

Wire segmentation in ultrasound images has always been a difficult problem due to the amount of noise in the B-scans. With the simple addition of a rubber membrane, we have shown how it is possible to solve this difficult segmentation problem. The wires can be segmented fully automatically and we are able to produce calibrations in real-time.

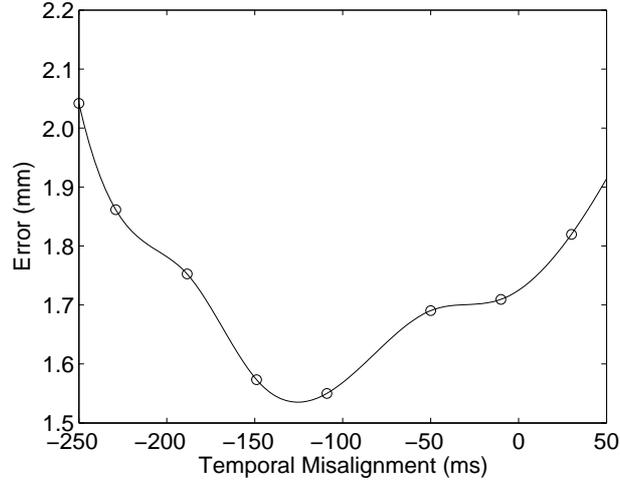


Figure 9: Mean 3D error of the fiducials when located in space using the position sensor readings at different temporal offsets.

We have noted that the estimation of the elevational offset is relatively poor when the calibration is performed using a Z-phantom. This estimation can be improved by capturing 8 images of the planar membrane. These images can be captured before calibration, and therefore the calibrations can still be produced in real-time.

In order to perform accurate spatial calibration using our Z-phantom, the user will need to first capture 8 images of the plane from an oblique angle. The phantom is then located in space by using a pointer. This process will take approximately 2 minutes. The phantom is then scanned and calibrations are produced in real-time. Re-calibration due to depth or zoom changes will therefore take a few seconds to complete.

By assuming that the spatial calibration is known, temporal calibration can be performed in real-time as well by moving the probe in a single direction. The temporal offset is found where the fiducials best match their locations in space. The low standard deviation of our temporal offset also compares favourably with the results achieved in the literature.

References

- J. D. Achenbach. *Wave Propagation in Elastic Solids*, volume 16 of *North-Holland series in applied mathematics and mechanics*. North-Holland, 1973.
- D. V. Amin, T. Kanade, B. Jaramaz, A. M. DiGioia III, C. Nikou, R. S. LaBarca, and J. E. Moody Jr. Calibration method for determining the physical location of the ultrasound image plane. In *Proceedings of the fourth International Conference on Medical Image Computing and Computer-Assisted Intervention, Lecture Notes in Computer Science*, volume 2208, pages 940–947. Springer-Verlag, 2001.
- K. S. Arun, T. S. Huang, and S. D. Blostein. Least-squares fitting of two 3-D point sets. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(5):698–700, 1987.
- L. G. Bouchet, S. L. Meeks, G. Goodchild, F. J. Bova, J. M. Buatti, and W. A. Friedman. Calibration of three-dimensional ultrasound images for image-guided radiation therapy. *Physics in Medicine and Biology*, 46:559–577, 2001.
- M. Burcher. *A force-based method for correcting deformation in ultrasound images of the breast*. PhD thesis, University of Oxford, Oxford, United Kingdom, 2002.
- R. M. Comeau, A. Fenster, and T. M. Peters. Integrated MR and ultrasound imaging for improved image guidance in neurosurgery. In *Proceedings of SPIE*, volume 3338, pages 747–754, 1998.
- S. Dandekar, Y. Li, J. Molloy, and J. Hossack. A phantom with reduced complexity for spatial 3-D ultrasound calibration. *Ultrasound in Medicine & Biology*, 31(8):1083–1093, 2005.
- P. R. Detmer, G. Bashein, T. Hodges, K. W. Beach, E. P. Filer, D. H. Burns, and D. E. Stradness Jr. 3D ultrasonic image feature localization based on magnetic scanhead tracking: in vitro calibration and validation. *Ultrasound in Medicine & Biology*, 20(9):923–936, 1994.
- A. Fenster, D. B. Downey, and H. N. Cardinal. Three-dimensional ultrasound imaging. *Physics in Medicine and Biology*, 46:R67–R99, 2001.
- A. H. Gee, R. W. Prager, G. H. Treece, and L. H. Berman. Engineering a freehand 3D ultrasound system. *Pattern Recognition Letters*, 24:757–777, 2003.
- A. H. Gee, N. E. Houghton, G. M. Treece, and R. W. Prager. A mechanical instrument for 3D ultrasound probe calibration. *Ultrasound in Medicine & Biology*, 31(4):505–518, 2005.
- M. J. Gooding, S. H. Kennedy, and J. A. Noble. Temporal calibration of freehand three-dimensional ultrasound using image alignment. *Ultrasound in Medicine & Biology*, 31(7):919–927, 2005.
- P-W. Hsu, R. W. Prager, A. H. Gee, and G. M. Treece. Rapid, easy and reliable calibration for freehand 3D ultrasound. *Ultrasound in Medicine & Biology*, 32(6):823–835, 2006.
- Q. H. Huang, Y. P. Zheng, M. H. Lu, and Z. R. Chi. Development of a portable 3D ultrasound imaging system for musculoskeletal tissues. *Ultrasonics*, 43:153–163, 2005.
- A. Khamene and F. Sauer. A novel phantom-less spatial and temporal ultrasound calibration method. In *Proceedings of the 8th International Conference on Medical Image Computing and Computer-Assisted Intervention, Lecture Notes in Computer Science*, volume 3750, pages 65–72. Springer-Verlag, 2005.
- T. Langø. *Ultrasound guided surgery: image processing and navigation*. PhD thesis, Norwegian University of Science and Technology, Trondheim, Norway, 2000.
- F. Lindseth, G. A. Tangen, T. Langø, and J. Bang. Probe calibration for freehand 3-D ultrasound. *Ultrasound in Medicine & Biology*, 29(11):1607–1623, 2003.

- S. Meairs, J. Beyer, and M. Hennerici. Reconstruction and visualization of irregularly sampled three- and four-dimensional ultrasound data for cerebrovascular applications. *Ultrasound in Medicine & Biology*, 26(2):263–272, 2000.
- L. Mercier, T. Langø, F. Lindsest, and D. L. Collins. A review of calibration techniques for freehand 3-D ultrasound systems. *Ultrasound in Medicine & Biology*, 31(4):449–471, 2005.
- D. M. Muratore and R. L. Galloway Jr. Beam calibration without a phantom for creating a 3-D freehand ultrasound system. *Ultrasound in Medicine & Biology*, 27(11):1557–1566, 2001.
- M. Nakamoto, Y. Sato, K. Nakada, Y. Nakajima, K. Konishi, M. Hashizume, and S. Tamura. A temporal calibration method for freehand 3D ultrasound system: a preliminary result. *International Congress Series*, 1256:1365, 2003.
- N. Pagoulatos, D. R. Haynor, and Y. Kim. A fast calibration method for 3-D tracking of ultrasound images using a spatial localizer. *Ultrasound in Medicine & Biology*, 27(9):1219–1229, 2001.
- R. W. Prager, R. N. Rohling, A. H. Gee, and L. Berman. Rapid calibration for 3-D freehand ultrasound. *Ultrasound in Medicine & Biology*, 24(6):855–869, 1998.
- R. W. Prager, A. H. Gee, and L. H. Berman. Stradx: Real-time acquisition and visualization of freehand three-dimensional ultrasound. *Medical Image Analysis*, 3(2):129–140, 1999.
- F. Rousseau, P. Hellier, and C. Barillot. Confhustus: A robust and fully automatic calibration method for 3D freehand ultrasound. *Medical Image Analysis*, 9(1):25–38, 2005.
- F. Rousseau, P. Hellier, and C. Barillot. A novel temporal calibration method for 3D ultrasound. *IEEE Transactions on Medical Imaging*, In Press, 2006.
- Y. Sato, M. Nakamoto, Y. Tamaki, T. Sasama, I. Sakita, Y. Nakajima, M. Monden, and S. Tamura. Image guidance of breast cancer surgery using 3-D ultrasound images and augmented reality visualization. *IEEE Transactions on Medical Imaging*, 17(5):681–693, 1998.
- G. M. Treece, A. H. Gee, R. W. Prager, C. J. C. Cash, and L. H. Berman. High-definition freehand 3-D ultrasound. *Ultrasound in Medicine & Biology*, 29(4):529–546, 2003.