3D Ultrasound Display using Optical Tracking

Peter H. Mills
Henry Fuchs

Department of Computer Science,
University of North Carolina,
Chapel Hill, N.C. 27599-3175 USA

Abstract

We present a method to capture a 3D ultrasound volume from a series of 2D cross-sectional images obtained with a conventional B-mode scanner. We use *optical tracking* of the ultrasound transducer to locate the digitized sections within the volume. In this paper we describe a simple optical tracking technique which uses two conventional video cameras and stereopsis, and discuss its freedom-of-motion effects on reconstruction and visualization. We are in the process of investigating its utility in an experiment imaging a human fetus. This technique is but a pilot forerunner of optical methods to be employed in a project whose goal is to achieve real-time interactive display of 3D ultrasound.

1. Introduction

1.1. 3D Ultrasound Imaging

Ultrasound imaging, in which reflected high-frequency sound waves are used to form pictures of internal anatomy, is a unique modality in that it is both non-destructive and can yield real-time display. Unlike other imaging modalities such as X-ray CT and MRI, however, conventional ultrasound *2D echographic* (2DE) scanners only produce 2D cross-sectional images, albeit in real-time. From this collection of 2D images it can be difficult to achieve comprehension of the three-dimensional (3D) structure and relationships of objects or organs [Gree82]. To effectively visualize anatomical characteristics, we want to have *3D acquisition and display* of ultrasound images; ideally, this visualization process should be real-time and interactive. In this paper we present a technique to capture and display a 3D ultrasound volume from a sequence of 2D cross-sectional images obtained with a standard linear scanner. We use *optical tracking* of the transducer to locate the digitized sections within the volume. In our pilot technique two video cameras view the position of a crown of lights attached to the transducer while the ultrasound scanning is performed: stereo triangulation then yields 6 degree-of-freedom tracking information. After reconstructing the 3D volume from the located planar samples, it can be displayed using volume-rendering and varifocal-mirror techniques. The notable feature of our technique is that it uses commercially available equipment for the optical tracking, so that anyone with access to similar equipment may employ it. We are in the process of applying our technique to generate a 3D image of a human fetus by post-processing a videotape of 2D scans.

1.2. Previous Work

The basic approach of generating a 3D image from a collection of 2D images with tracked position and orientation has been applied many times in the field of ultrasound imaging. In common with all these methods are these concerns:

- **Acquisition** of 2DE cross-sections.
  - **Tracking** the transducer arm to locate each slice in the volume.
  - **Digitizing** the 2D ultrasound image.
  - **Synchronizing** the 2DE images with the tracker location.
- **Visualization** of 3D objects.
  - **Segmentation** of the image into objects.
  - **Reconstruction** of the 3D volume or surface.
  - **Rendering** the picture.

The tracking methods differ both in the equipment used and in the *degree-of-freedom* (DOF), or *dimensionality*, of the information produced. For example, a 6-DOF tracker will yield 3-dimensional (x, y, z) *position* information, as well as 3 angles describing *orientation*. Brinkley et al. [Brin78] employed 6-DOF acoustic tracking, using spark gaps as sound sources, to locate the transducer. Organ contours were then manually traced from the 2DE images, and a finite-element-mesh, or polygonal skin, was reconstructed from the positioned contours. Polygonal mesh reconstructions, or *lattices*, were also generated from the traced borders of stop-frame images by Ghosh et al. [Ghos82], who mechanically tracked a 1-DOF stationary rotating sector scanner. Stickles and Wann [Stic84] used a 3-DOF mechanical arm to locate manually sampled data points, and then reconstructed a mesh from the samples. Nikravesh et al. [Nikr84] employed a 6-DOF mechanical arm for tracking, then also produced a lattice.

More recent implementations reconstruct the image volume directly from the planar samples before surface formation, if any, is done. McCann et al. [McCa88] use a 1-DOF rotation method similar to that of Ghosh [Ghos82], but then reconstruct the image volume and display it using tools from the ANALYZE II system [Robb87]. Lalouche et al. acquire a set of regularly-spaced parallel 2D image slices via a 1-DOF serial imager, then interpolate the image and use Pixar volume-rendering display routines [Dreb88]. In a hybrid reconstruction technique, Raichlen et al. [Raic86] employ a 1-DOF reorientable mechanical slider to obtain parallel slanted cross-sections at regular intervals, manually segment the 2D images into objects, and then reconstruct and display the volume using the *cuberille* techniques of G. Herman and J. Udupa. Note that in all of the cases described above, the images were reconstructed and visualized in *post-processing* steps, rather than in *real-time*.

These latter techniques, which perform volume reconstruction, *constrain the tracking* degree-of-freedom to be less than 6-DOF in order to exploit the parallel geometry to generate uniformly sampled images.

1.3. Optical Tracking: Pilot Technique

We focus here on the issue of tracking with 6-DOF for 3D ultrasound capture and display. We advance three main propositions.

- We advocate using 6-DOF optical tracking.
- Conventional technology, or off-the-shelf components, should be used if possible.
- The effects of the freedom of motion on volume reconstruction from 2D oblique slices must be addressed.

As a first attempt in using optical tracking to acquire 3D ultrasound, we applied a simple optical tracking technique in an experiment to capture a 3D volume of a human fetus. To date we have recorded the incrementally acquired 2D slices on videotape, and are now in the progress of post-processing the results.

The optical tracking of the transducer, the business end of the ultrasound scanner, was achieved using simple stereometric techniques. An antler holding four incandescent lights was attached to the transducer. Two video cameras were mounted in the corners of the room to view the lights' position; their video output, as well as that of ultrasound scanner, was recorded on videotape. In order to know which video frames were coincident, the tapes were synchronized by using a single time-code-generator to record event times on the audio tracks of all three tapes. To locate the positions and focus of the cameras, a view of a cube with an LED at each corner sufficed. Calibration of the location of the antlers relative to the attached transducer was also required: we first physically measured the relative position, but future experiments will use calibration by scanning a known phantom object.

Once the data was acquired, reconstruction of the 3D volume from the oblique cross-sections proceeds by digitizing the videotapes, computing the position of each ultrasound slice from the simultaneous views of the LEDs, and then deriving each voxel value from an average of the sampled cross-sections. The volume, after some manual editing, can be viewed on a conventional raster display using Marc Levoy's volume rendering algorithm [Levo88]. The volume can also be viewed on the true-3D varifocal-mirror display system developed here at UNC [Fuch82]. Figure 1 gives an overview of the entire process.

1.4. Real-time Visualization

The experiment just described dealt with the generation of 3D ultrasound images using conventional technology. It also served to provide pilot data for a larger project whose goal is to achieve *real-time* 3D acquisition and display of ultrasound data. An interim system for real-time display of 3D structures from a series of incrementally acquired 2D ultrasound slices is being developed by H. Fuchs and R. Ohbuchi [Ohbu90]. This system will achieve real-time interactivity by incrementally updating the reconstructed and displayed 3D ultrasound volume as each 2D slice comes in from the scanner. A 3-DOF mechanical arm yielding (x, y) position and z-axis rotation values currently provides the tracking mechanism.

These are steps towards the ultimate goal of a real-time system to acquire and display 3D ultrasound images. This future system will use a true 3D echographic scanner being developed by Dr. Olaf von Ramm and colleagues [Shat84]. The parallel 3D acquisition device will be coupled to a rendering machine and perhaps eventually to a headmounted display, wherein the physician can wear a visor giving a real-time stereo perspective view of the ultrasound data overlaid onto his view of the patient. This system is also envisioned to rely on real-time optical trackers for both the ultrasound transducer and the user's head-mounted display.

In the sections that follow we describe our interim

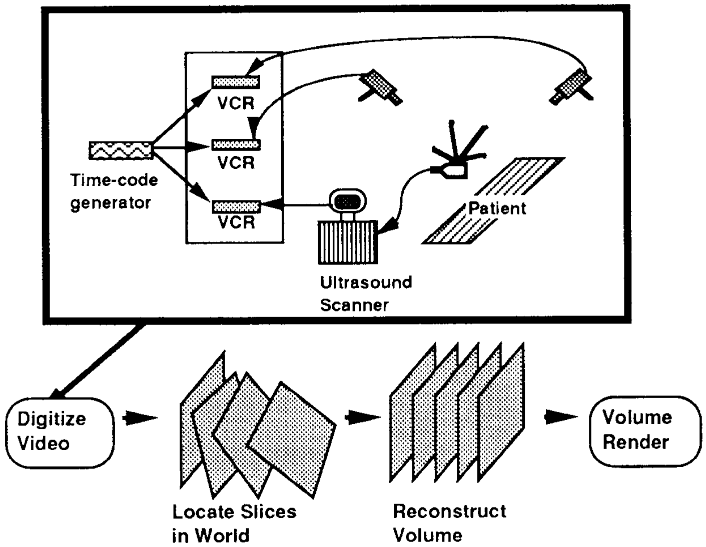

Figure 1: 3D ultrasound display using optical tracking

pilot optical tracking technique, the proposed reconstruction algorithm and issues, and the rendering tools for display.

2. Incremental Acquisition

Capturing a 3D volume by incrementally acquiring 2D ultrasound slices requires *tracking* the transducer, *synchronizing* the 2DE images with the tracked locations, and having *calibrated* the 2DE image plane relative to the tracked object. Before describing how our experimental setup deals with each of these issues, we briefly discuss optical tracking in general.

2.1. Tracking Methods

There are several methods of determining the position and orientation of the transducer, or any object, in real-time. Among them are:

- **Mechanical** trackers (e.g., robot arms).
- **Magnetic** trackers (e.g., Polhemus [Raab79]).
- **Acoustic** trackers.
- **Inertial motion** trackers (accelerometers).
- **Optical** trackers.

Optical tracking systems typically use two or more cameras to view a set of reference lights, or *beacons*, attached to the object to be tracked. Given the observed 2D image positions of the beacons and the known fixed 3D positions of the cameras, stereo triangulation can be used to compute the positions of the beacons, and hence of the tracked object, in space. In the dual to this case, *inside-out* tracking, camera sensors are placed on the tracked object and the beacons are fixed in location [Wang90b]. Another refinement to the optical tracking paradigm discards the use of artificial beacons entirely: these *natural environment* trackers sense changing light patterns in the environment as a basis for the tracking process. An example of this is Bishop's *Self-tracker* [Bish84]. Wang implemented a prototype inside-out tracker using off-the-shelf components, in which lateral-effect photodiodes view ceiling-mounted LED beacons [Wang90a].

Optical tracking techniques have the advantages that they do not greatly restrict the motion of the user, are relatively insensitive to environmentally induced distortion, have a large working volume, and can be made fast and accurate [Wang90b]. Wang's technique achieved, for a small field of view, an accuracy of roughly 1 part in 1000, or 2mm of translation resolution and up to .1 degree of rotational resolution. For a summary of other tracking methods, see [Wang90b].

2.2. Optical Tracking Method

As a first trial for optical tracking in 3D ultrasound we wanted a simple technique that required only commercially available equipment. A tradeoff was made between possibly low tracking resolution and the ease of pilot implementation. The implementation details follow.

Experimental Setup. In our experiment, as shown in Figure 1, we mounted two video cameras in corners of the room. A set of four prongs, an incandescent light at the tip of each, is attached to the transducer of the ultrasound machine, in our case an Ultramark4 Ultrasound system. Three VCR's recorded the 2DE ultrasound data and each camera's view for later post-processing. We synchronized coincidental frames of the three videotapes by placing on the audio track of each videotape time-code signals from an Evertz time-code generator. Figures 2 and 3 show one recorded frame from each of the cameras.

Digitization. To digitize synchronous video frames from each of the two cameras and the 2DE image, we employed a Lyon-Lamb VAS-IV animation controller coupled to a Harris 642 freeze-frame, whose output was routed to a digitizer. For each time-code (HH-MMSSFF), the Evertz time-code reader established the corrspondence to frames in the videotape, and the animation controller found the desired frame and commanded the frame-synchronizer to freeze the frame, after which it was digitized.

Locating 2D slices. The method of determining the location of each 2D ultrasound slice in the world is shown in Figure 5. For each video frame of 2D ultrasound data, we digitize the corresponding two simultaneous camera views, as shown in Figure 4. We then find the beacons in the 2D camera images by simple image processing. Spurious beacons from reflections are discarded. The position and orientation of the antler in space is then derived using stereometric techniques described by Sutherland [Suth74]. However, this last step presupposes we know which pairs of lights, one from each camera view, correspond to the same beacon. To establish this point correspondence, we test all permutations of pairing lights from each camera view. For each permutation, we derive the location of all beacons, and see if this set matches the known configuration of the antlers. This technique requires an asymmetric pattern of beacons, and is somewhat similar to that used by Ganapathy for 6D optical tracking using a single camera [Gana85].

This yields the location of the antlers in space for that 2DE image. To determine the location and physical size of the 2DE image in the world requires knowing the relative location of the transducer's image relative to the tracked object, i.e. the antlers.

2.3. Calibration

We designed a simple method to determine the aforementioned relation of the 2D ultrasound image and the tracked antlers. For this we require a phantom calibration object of known size, as shown in Figure 6. This phantom is similar to that described by Joynt [Joyn82], and consists of nylon strings strung between plexiglass plates: the whole structure is immersed in a water-filled tank. For any 2DE image of the phantom, from the intersected positions of the strings in the 2DE image we can derive the physical size and location of the 2DE image relative to the phantom. Then given the rigid-body motion of the joined transducer and tracked beacons, two sampled positions of the join give T_t as the motion transform of the beacons as tracked by the cameras, T_i as the motion transform of the 2DE image as tracked by the phantom, and T_e as the desired embedding of the 2DE image into beacon space. Linear least-squares fit yields T_e.

3. Reconstruction

To reconstruct the 3D volume from the oblique 2DE slices, the voxel (volume-element) values are derived from an average of the resampled oblique cross-sections. We developed what is in some ways a variant of Ohbuchi's *incremental reconstruction* [Ohbu90], extended to handle 6 degrees of freedom.

As with other techniques for volume reconstruction from slices [Raic86, McCa88], the three basic steps involved in reconstruction are:

- **Resampling**, or injecting, the oblique planes into the volume by distributing contributions to nearby voxels.
- **Interpolating**, or hole-filling, those voxels outside an injected plane.
- **Smoothing** the resultant volume.

The unfilled voxels are typically filled by linearly interpolating along a major axis between temporally adjacent 2DE slices.

While the optical tracking setup was reasonably straightforward, its use introduces some subtleties into the volume reconstruction that make it more difficult than that from parallel cross-sections or linear sweeps. First, inherent in 3D image reconstruction from incremental 2DE cross-sections is the temporal delay between slices which poses problems due to patient and organ motion or change. Secondly, the nature of optical tracking allows a greater degree of freedom of movement with the transducer than, for example, a robot-arm tracking method. This freedom of motion results in slice samples often at highly oblique angles to each other, instead of a nearly co-planar sweep covering the volume. During reconstruction, these temporally removed oblique passes contribute to the voxel values.

Indeed, even as a post-processing step, volume reconstruction from optically tracked 2DE slices poses many of the same aging problems described by R. Ohbuchi for his *incremental reconstruction* [Ohbu90]. There, the

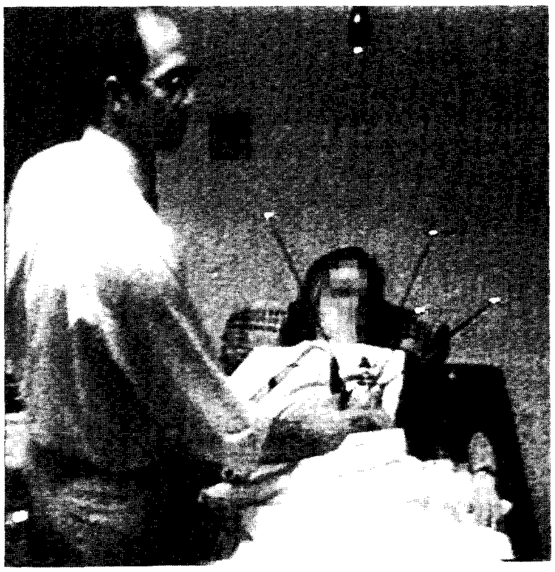

Figure 2: Foot camera view

Figure 3: Side camera view

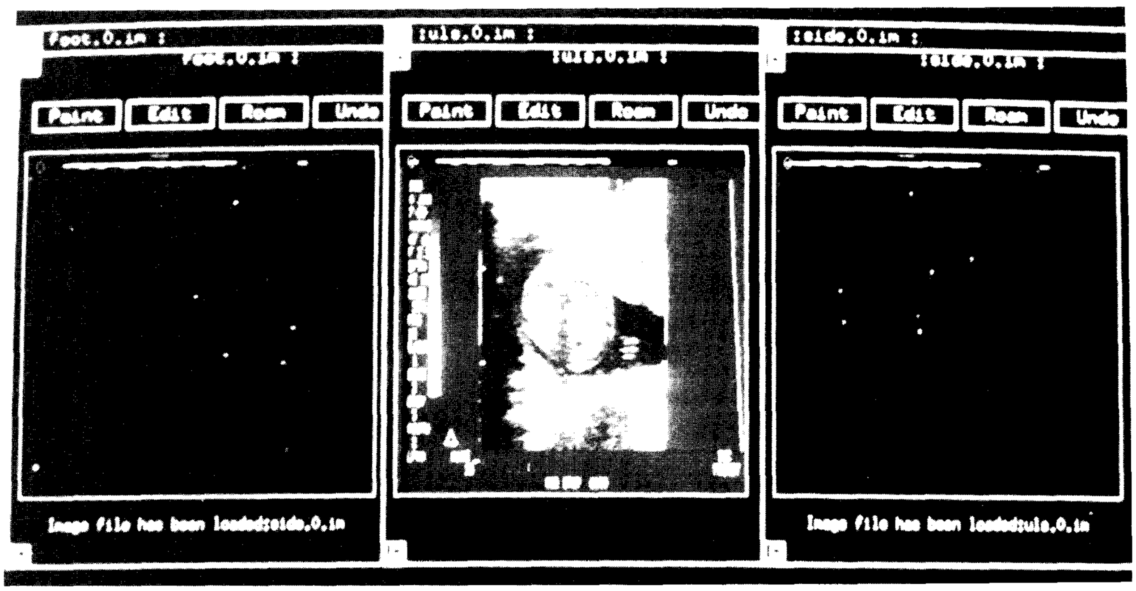

Figure 4: Synchronous camera views for ultrasound slice

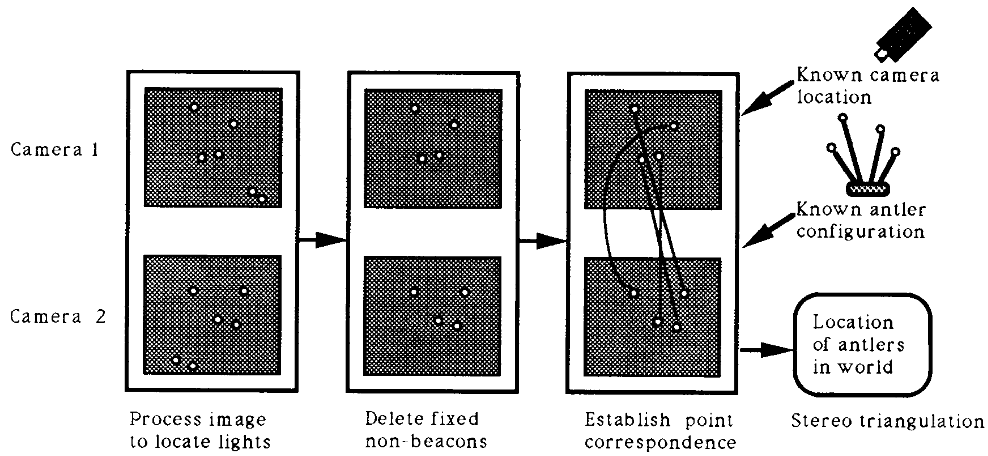

Figure 5: Optical tracking using stereo cameras

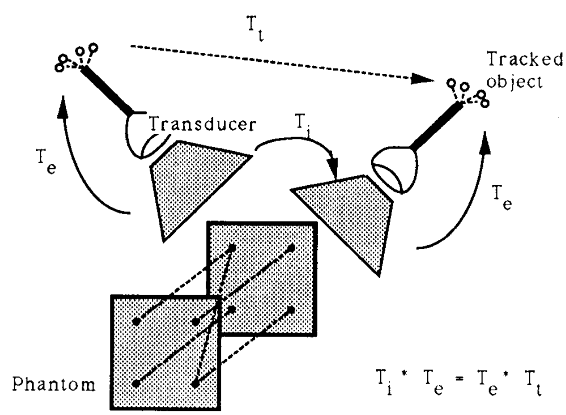

Figure 6: Calibration

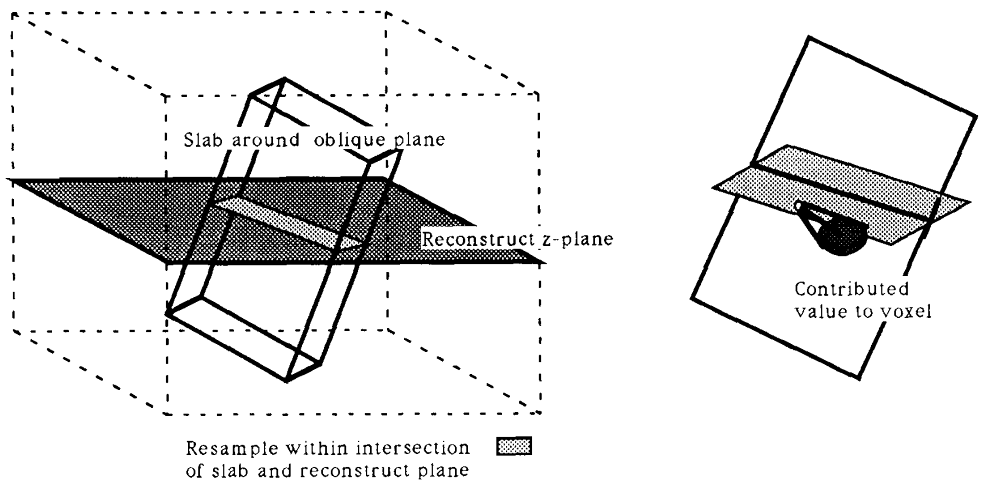

Figure 7: Resampling for voxels near oblique 2DE plane

volume is reconstructed incrementally as each new 2DE slice is acquired. Ohbuchi resamples the new 2DE plane and fills empty voxels between the new and previous slice by *forward mapping*, or *splatting*, contributions from linearly-interpolated input 2DE values to neighboring voxels. The problem of combining new contributions with values from older slices is attacked by the *decay method* and the *motion vector method.* Unfortunately, for 6-DOF tracking the motion vector method fails since we are no longer necessarily sweeping back and forth. Furthermore, for 6-DOF his weight buffer must be extended from 2D to 3D, since there is more than one axis of rotation.

In our reconstruction method we also perform sampling by weighted distribution of 2DE values, keeping at each voxel $(\sum weight_i, \sum weight_i * value_i)$. [To normalize this value merely divide through by the sum of weights]. However, instead of forward mapping, we use the dual method of *backward mapping* as shown in Figure 7. Those voxels in the intersection of a given reconstruction plane with a slab around the 2DE plane have contributions added. Interpolation between slices then operates in voxel space, sacrificing accuracy for speed of computation. We are still investigating possible reconstruction techniques to reduce the local effects of temporally distant oblique slices.

4. Visualization

Since we restricted our experiment to post-processing, standard volume-rendering techniques [Levo88] suffice for rendering. Volume-rendering methods directly picture the volume of data without first fitting a surface defined by geometric primitives. Although no segmentation or surface detection is needed, some coarse manual editing to remove obscuring tissue or artifacts is anticipated.

5. Summary and Future Work

We described methods for the capture and display of 3D ultrasound from a series of 2DE images whose location in the volume is tracked optically. The pilot system uses off-the-shelf components for the tracking mechanism. An experiment using this pilot system to image a human fetus will soon be completed.

As a next step in optical tracking we plan to repeat the experiment described here, replacing the twin video-camera tracking with a custom optical tracking wand being constructed at UNC which can output position and orientation in real-time [Wang90a]. Finally, we plan to incorporate this new optical tracker into the *incremental* real-time 3D ultrasound system under development by Ohbuchi and Fuchs at UNC. The *incremental reconstruction* algorithms of this system need to be adapted to accomodate the greater degree of freedom in optical tracking. These accomodations include interpolation functions for forward mapping between oblique slices, as well as techniques for properly weighting older slices.

Acknowledgements

The authors wish to thank Vern Katz, M.D. in the Department of Obstectrics and Gynecology at the UNC School of Medicine, for his assistance with the ultrasound exam. We also wish to thank Robert Strain for the photography. This research was supported in part by the National Science Foundation under grant number CDR-86-22201.

References

[Bish84] Bishop, T.G. and Fuchs, H., "The self-tracker: A smart optical sensor on silicon", in *Proceedings of Conference on Advanced Research in VLSI*, (MIT Press, 1984).

[Brin78] Brinkley, J.F., Moritz, W.E. and Baker, D.W., "Ultrasonic three-dimensional imaging and volume from a series of arbitrary sector scans", *Ultrasound in Med. & Bio.*, Vol.4 (1978) pp.317-327.

[Dreb88] Drebin, R.A., Carpenter, L. and Hanrahan, P., "Volume Rendering", *Computer Graphics*, SIGGRAPH'88 Conference Proceedings, Vol.22, No.4 (August 1988) pp. 65-74.

[Fuch82] Fuchs, H., Pizer, S.M., et al., "Adding a True 3-D Display to a Raster Graphics System", *IEEE Computer Graphics and Applications*, Vol. 2 (1982) pp.73-78.

[Gana85] Ganapathy, S.K., "Real-Time Motion Tracking Using a Single Camera", AT&T Bell Labs TM 11358 (1985).

[Ghos82] Ghosh, A., Nanda, C.N. and Maurer, G., "Three-dimensional reconstruction of echo-cardiographic images using the rotation method", *Ultrasound in Med. & Biol.*, Vol.8, No.6 (1982) pp.655-661.

[Gree82] Greeenleaf, J.F., "Three-Dimensional Imaging in Ultrasound", *Journal of Medical Systems*, Vol.6, No.6 (1982) pp.579-589.

[Joyn82] Joynt, L. and Popp, R.L., "The concept of three dimensional resolution in echocardiographic imaging", *Ultrasound in Med.& Biol.*, Vol.8, No.3 (1982) pp.237-247.

[Lalo89] Lalouche, R.C. et al., "Three-dimensional reconstruction of ultrasound images", in: *Medical Imaging III: Image Processing*, Roger H. Schneider, Samuel J. Dwyer III, R. Gilbert Jost (eds), Proc. SPIE 1092 (1989), pp.59-66.

[Levo88] Levoy, M., "Display of Surfaces from Volume Data", *IEEE Computer Graphics and Applications*, Vol.8, No.5 (May 1988) pp.29-37.

[McCa88] McCann, H.A. et al., "Multidimensional Ultrasonic Imaging for Cardiology", *Proc. of IEEE*, Vol.76, No.9 (Sept.1988) pp.1063-1073.

[Nikr84] Nikravesh, P.E. et al., "Computerized three-dimensional finite element reconstruction of the left ventricle from cross-sectional echocardiograms", *Ultrasonic Imaging*, Vol.6 (1984), pp.48-59.

[Ohbu90] Ohbuchi, R. and Fuchs, H., "Incremental 3D Ultrasound Imaging from a 2D Scanner", *Proc. of the 1st Conference on Visualization in Biomedical Computing*, (Atlanta, GA, May.1990), in these proceedings.

[Raab79] Raab, F.H. et al., "Magnetic Position and Orientation Tracking System", *IEEE Transactions on Aerospace and Electronic Systems*, Vol.15 (1979) pp.709-718.

[Raic86] Raichlen, J.S. et al., "Dynamic Three-Dimensional Reconstruction of the Left Ventricle From Two-Dimensional Echocardiograms", *Journal. Amer. Coll. of Cardiology*, Vol.8, No.2 (Aug.1986), pp.364-370.

[Robb87] Robb, R.A., "A workstation for interactive display and analysis of multi-dimensional biomedical images", in: *Computer Assisted Radiology*, Proceedings of the International Symposium CAR'87, H.U. Lemke et al., eds. (Springer-Verlag, Berlin, 1987) pp. 642-656.

[Shat84] Shattuck, D.P. et al., "Explososcan: A Parallel Processing Technique for High Speed Ultrasound Imaging with Linear Phased Arrays", *JASA*, 75(4) (1984) pp.1273-1282.

[Stic84] Stickels, K.R. and Wann, L.S., "An Analysis Of Three-Dimensional Reconstructive Echocardiography", *Ultrasound in Med. & Biol.*, Vol.10, No.5 (1984) pp.575-580.

[Suth74] Sutherland, I.E., "Three-Dimensional Data Input by Tablet", *Proc. of IEEE*, Vol.62, No.4 (April.1974) pp.453-461.

[Wang90a] Wang, Jih-Fang, Chi, V. and Fuchs, H., "A Real-time Optical 3D Tracker for Head-mounted Display Systems", to appear in: *Proceedings of the 1990 Symposium on Interactive 3D Graphics*, (Snowbird, Utah, March.1990).

[Wang90b] Wang, Jih-Fang, "A Real-time Optical 6D Tracker for Head-mounted Display Systems", Ph.D. Thesis, Dept. of Computer Science, Univ. of North Carolina, Chapel Hill, N.C. (1990).