Engineering Tripos Part IIB & EIST Part II FOURTH YEAR

Module 4F9: Medical Imaging & 3D Computer Graphics

Solutions to 2003 Tripos Paper

1. *Ultrasound wave propagation*

(a) $k_1 = \omega/c_1$. [10%]

(b) The pressure on both sides of the boundary must be equal:

$$\begin{aligned} p_i|_{x=0} + p_r|_{x=0} &= p_t|_{x=0} \\ \Rightarrow A_1 + B_1 &= A_2 \end{aligned}$$

This gives us one equation relating A_1, A_2 and B_1. [10%]

(c) The particle velocities, normal to the interface, must be equal at the boundary.

$$u_i|_{x=0} + u_r|_{x=0} = u_t|_{x=0}$$

The linear inviscid force equation is:

$$\begin{aligned} \rho\frac{\partial u}{\partial t} &= -\frac{\partial p}{\partial x} \\ \Rightarrow u &= -\left(\frac{1}{\rho}\right)\int\left(\frac{\partial p}{\partial x}\right)dt \end{aligned}$$

now if $p(x,t) = p(x)e^{i\omega t}$ then

$$u = -\left(\frac{1}{i\omega\rho}\right)\frac{\partial p}{\partial x}$$

and if we substitute the expressions for p given in the question we get:

$$\frac{A_1}{\rho_1 c_1} - \frac{B_1}{\rho_1 c_1} = \frac{A_2}{\rho_2 c_2}$$

This gives us a second relationship between A_1, A_2 and B_1. [30%]

(d)

$$\begin{aligned} R &= \frac{B_1^2}{A_1^2} \\ T &= \frac{A_2^2\rho_1 c_1}{A_1^2\rho 2 c_2} \end{aligned}$$

[20%]

(e) Solving the two relationships between A_1, A_2 and B_1 together we can obtain expressions for B_1/A_1 and A_2/A_1 in terms of ρ_1, ρ_2, c_1 and c_2. This leads to

$$\begin{aligned} R &= \left(\frac{\rho_2 c_2 - \rho_1 c_1}{\rho_2 c_2 + \rho_1 c_1}\right)^2 \\ T &= \frac{4\rho_1 c_1 \rho_2 c_2}{(\rho_1 c_1 + \rho_2 c_2)^2} \end{aligned}$$

[30%]

Assessor's comments: *(Attempted by 31/42 candidates, average mark 11.2/20). This question covered the bookwork describing the way an ultrasound wave crosses an interface between two materials. Some candidates had a full understanding of the material and gained high marks, but several were not able to remember the basic definitions taught in the lectures.*

2. **Radon and X-ray transforms**

(a) The three-dimensional radon transform maps a function $f(x, y, z)$ to the set of its integrals over planes perpendicular to unit vectors $\boldsymbol{u}$ at a distance s from the origin.

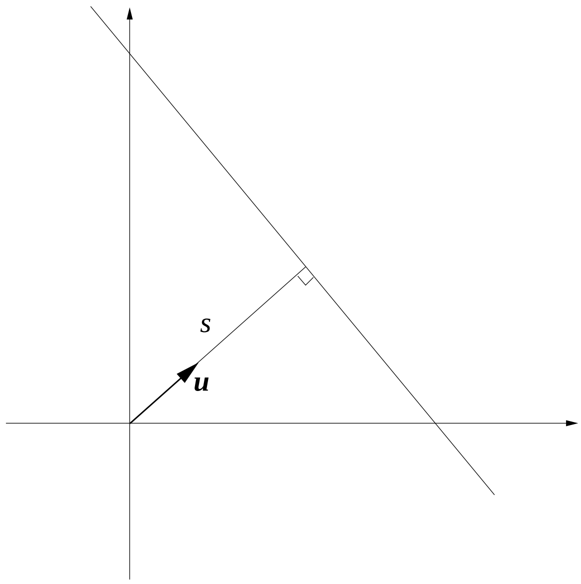

$$\mathbf{R}f(s, \boldsymbol{u}) = \int_{\boldsymbol{r}.\boldsymbol{u}=s} f(\boldsymbol{r}) d\boldsymbol{A}$$

where $d\boldsymbol{A}$ is a small element of the plane perpendicular to $\boldsymbol{u}$ a distance s from the origin, i.e. the plane defined by $\boldsymbol{r}$ when $\boldsymbol{r}.\boldsymbol{u} = s$.

The X-ray transform maps a function onto the set of its line integrals. Consider a line passing through a point with position vector $\boldsymbol{x}$, with direction $\boldsymbol{v}$.

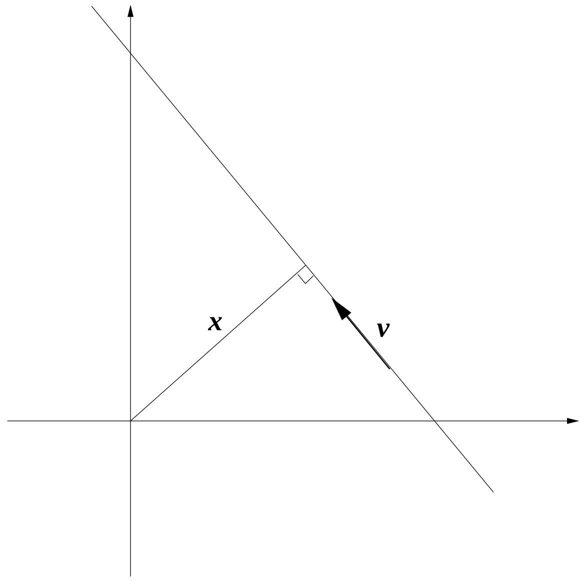

$$\mathbf{P}f(\boldsymbol{x}, \boldsymbol{v}) = \int\limits_{\boldsymbol{r}=\boldsymbol{x}+\lambda\boldsymbol{v}} f(\boldsymbol{r})|d\boldsymbol{r}|$$

Note: in any individual projection, using the X-ray transform, the direction of $\boldsymbol{v}$ is fixed, and $\boldsymbol{x}$ varies such that $\boldsymbol{x}.\boldsymbol{v} = 0$. [20%]

(b) The X-ray transform is an integral along a line perpendicular distance s from the origin. The function is spherically symmetric so the X-ray transform will not depend on the direction of the line, it will only be a function of s. Let r be the distance along the line. The X-ray transform therefore equals

$$\begin{aligned}
\mathbf{P}f(s) &= \int\limits_{-\infty}^{+\infty} f \mathrm{d}r \\
&= \int\limits_{-\infty}^{+\infty} \frac{1}{(x^2+y^2+z^2)^2} \mathrm{d}r \\
&= \int\limits_{-\infty}^{+\infty} \frac{1}{(s^2+r^2)^2} \mathrm{d}r \\
&= \frac{\pi}{2s^3}
\end{aligned}$$

[40%]

(c) The radon transform is an integral over a plane perpendicular distance s from the origin. Let p and q be distances in perpendicular directions in the plane of the integral. The radon transform therefore equals

$$\begin{aligned}
\mathbf{R}f(s) &= \int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} f \mathrm{d}q\mathrm{d}p \\
&= \int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} \frac{1}{(x^2+y^2+z^2)^2}\mathrm{d}q\mathrm{d}p \\
&= \int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} \frac{1}{(s^2+p^2+q^2)^2}\mathrm{d}q\mathrm{d}p \\
&= \frac{\pi}{2}\int_{-\infty}^{+\infty} \frac{1}{(s^2+p^2)^{\frac{3}{2}}}\mathrm{d}p \\
&= \frac{\pi}{s^2}
\end{aligned}$$

[40%]

Assessor's comments: *(Attempted by 10/42 candidates, average mark 10.3/20) This question tested the candidates' understanding of the X-ray and radon transforms. Most candidates produced plausible definitions of the X-ray and radon transforms. In the algebraic part of the question, all the candidates found the solution of the line and surface integrals difficult. Some were not even able to formulate the expressions for the integrals correctly.*

3. **Marching Cubes and Tetrahedra**

(a) The volume of a tetrahedra is given by 1/3 base area times height. Each cube thus contributes a volume of 1/48. There are 8 cubes so we have a total volume of $8 \times 1/48 = 1/6$. [25%]

(b) By symmetry, the solid edges form a cubic lattice. We therefore can consider the following slice:

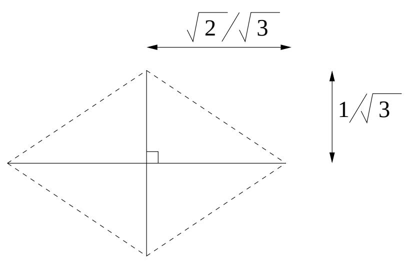

The length of the dashed sides is thus one unit. [30%]

(c) By counting from the diagram we see that there are 24 segments. [15%]

(d) Each little tetrahedron, with side lengths of $1/2$ and $1/\sqrt{3}$, that contributes to the volume has a base area of $1/(6\sqrt{2})$ and a height of $1/\sqrt{6}$. Its volume is thus $1/(36\sqrt{3})$. There are 24 such segments, giving a total volume of $2/(3\sqrt{3}) = 0.385$. [30%]

Assessor's comments: *(Attempted by 29/42 candidates, average mark 10.8/20). This question covered the construction of 3D surfaces using the marching cubes and marching tetrahedra algorithms. Candidates were generally able to solve the marching cubes problem in part (a). They found the geometry in parts (b) and (c) surprisingly difficult because it required careful visualisation in 3D and an awareness of symmetry.*

4. **Shape Based Interpolation**

(a) The diagram below shows the results of linear interpolation with a thick line at the boundary.

2	5	8	7	8	10	8	6	4	1
2	4.5	7	6	7	9	8	6.5	4.5	1.5
2	4	6	5	6	8	8	7	5	2
2	3.5	5	4	5	7	8	6.5	5.5	2.5
2	3	4	3	4	6	8	8	6	3

[25%]

(b) The diagram below shows the result of one dimensional shape based interpolation with a thick line at the boundary.

−1	0	1	2	3	2	1	0	−1	−2
	−1	0	1	2				−3/4	
	−2	−1	0	1				−1/2	
	−3	−2	−1	0				−1/4	
−5	−4	−3	−2	−1	0	1	1	0	−1

[35%]

(c) First we set up distance fields in each of the two slices that are provided.

−3	−2	−2	−1	−1	−1	−1	−2	−3	−4
−2	−1	−1	0	0	0	0	−1	−2	−3
−1	0	0	1	1	1	0	0	−1	−2
−2	−1	0	0	0	0	−1	−1	−2	−3
−3	−2	−1	−1	−1	−1	−2	−2	−3	−4

−5	−4	−3	−2	−1	−1	−1	−1	−1	−2
−4	−3	−2	−1	0	0	0	0	0	−1
−5	−4	−3	−2	−1	0	1	1	0	−1
−5	−4	−3	−2	−1	0	0	0	−1	−2
−6	−5	−4	−3	−2	−1	−1	−1	−2	−3

We then interpolate the distance fields to create a slice half way between the given planes of data.

−	−	−	−	−	−	−	−	−	−
−	−	−	−0.5	0	0	0	−0.5	−1	−
−	−2	−1.5	−0.5	0	0.5	0.5	1	−0.5	−
−	−	−1.5	−1	−0.5	0	−0.5	−0.5	−	−
−	−	−	−	−	−	−	−	−	−

The object is taken as every point with a value of zero or above.

This answer has been based on a sequence of the form (-3, -2, -1, 0, 1, 2, 3), with the threshold just below zero. Other sequences, for instance (-20, -15, -10, -5, +5, +10, +15, +20), resulting in slightly different answers, are also acceptable. [40%]

Assessor's comments: *(Attempted by 12/42 candidates, average mark 10.8/20). The first part of this question tested candidates' understanding of shape-based interpolation in 2D. The later part of the question addressed similar issues in 3D. Although unpopular, this question resulted in a wide range of marks as those who understood the principle behind the algorithm achieved high scores, whereas those with a muddled understanding collected few marks.*

5. **Volume Rendering**

(a) Volume rendering is the term used to describe the direct visualisation of a 3D voxel array without explicit segmentation into surfaces. Rays of light (usually parallel) are cast through the volume where they are shaded and attenuated by the individual voxels. The rays exiting the volume impinge on an image plane to create the rendering, in a manner analogous to taking an X-ray. The colour and opacity of the voxels are set using classification techniques, often heuristic in nature. The volume may be re-oriented with respect to the image plane to construct views from different viewpoints, clipping planes can be used to look behind occluding objects, and the colours and opacities can be adjusted interactively to emphasize different aspects of the volumetric data.

Compared with surface rendering, the advantages of volume rendering for volume visualisation are:

- No need to explicitly extract surfaces — the difficult segmentation problem is avoided.
- Ability to visualise more of the data set at a time: animation helps to disambiguate the renderings.
- Ability to interactively view different aspects of the volume by adjusting the clipping planes and the material parameters.

The disadvantages of volume rendering are:

- A lot of parameters need to be set by trial and error (this can be construed more positively as beneficial user interaction).
- Considerable computational expense each time you render.
- No geometric properties e.g. volume / surface area.

[25%]

(b) (i) Minimum intensity projection will show the vasculature. (ii) Maximum intensity projection will show the cartilage. [20%]

(c) The area of box B is $80 \times 80 = 6400\text{mm}^2$. To find the area of box A, the length of the box is $80\sqrt{2}$, and the width is $40/\sqrt{2} + 80 - 40\sqrt{2}$. Thus the area is 5851mm^2. So we choose box A.

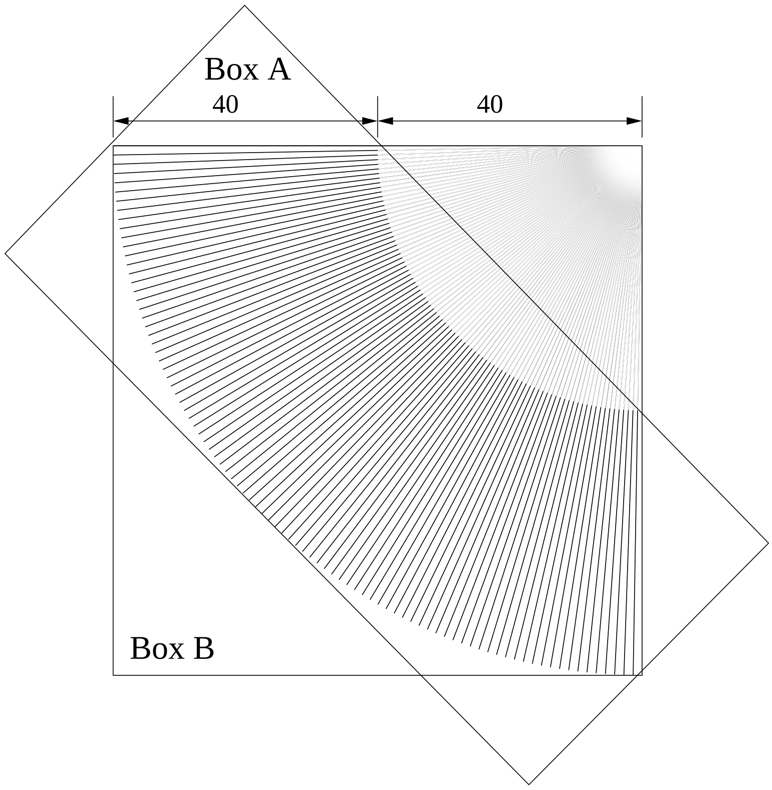

For no loss of information when resampling at 45° we require voxels $0.1/\sqrt{2}$ mm cubed. We therefore need a voxel array $1600 \times 732 \times 425 = 497,760,000$, or about 475 megabytes of memory, at one byte per voxel. [30%]

(d) The memory required for the raw data is $91 \times 300 \times 400 = 10920000$ bytes. The required ratio is thus $497760000/10920000 = 45.6$ [10%]

(e) Nearest neighbour, neighbourhood interpolation with a 3D Gaussian kernel, neighbourhood interpolation with a 3D kernel shape based on the local shape of the ultrasound resolution cell, radial basis functions, b-spline interpolation. [15%]

Assessor's comments: *(Attempted by 22/42 candidates, average mark 11.1/20). This question was about volume rendering, and the practicalities of using traditional volume rendering algorithms to visualise 3D ultrasound data. Several candidates ignored the wording of the question and discussed volume rendering exclusively in the context of CT data. The numerical question about resampling inhomogeneous data onto a regular voxel array was generally well done, and most candidates were able to suggest interpolation algorithms as required in part (e).*

6. **Shadow Rendering**

(a) The order of application is: first modelview matrix, then the projection matrix then clipping. The modelview matrix determines the orientation and position of a line with respect to the rest of the drawing, the project matrix determines the overall mode of viewing, usually either perspective projection or orthogonal projection, and the global scale. The clipping determines the size of the volume of space to be rendered, called the *view volume*. [20%]

(b)

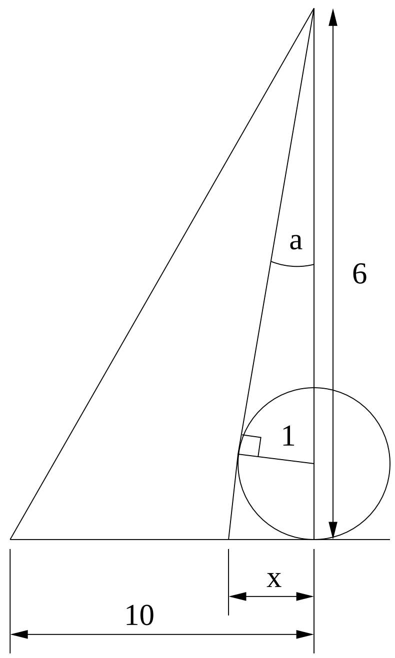

Angle $a = \sin^{-1}(1/5)$

Length $x = 6\tan(a) = 1.2247$

This covers $1.2247 \times 128/10 = 15.7$ pixels. So the width of the sphere is 31 pixels. [20%]

(c)(i)

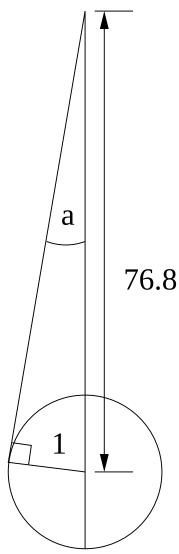

Distance from sphere to viewpoint $= \sqrt{50^2 + 50^2 + 30^2} = 76.8$

Angle $a = \sin^{-1}(1/76.8)$

Pixels along sphere radius $= 6\tan(a) \times 128/10 = 1.0$

So the width of the sphere is 2 pixels in the shadow z-buffer. [15%]

(ii)

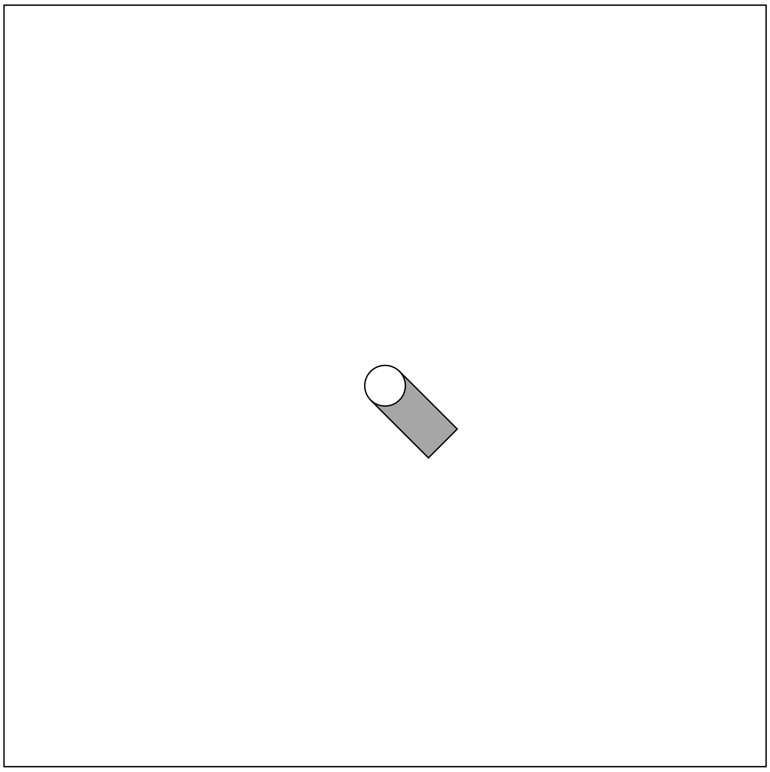

[15%]

(iii) The result depends on the positions of the near and far clipping planes used for the shadow z-buffer. If the clipping planes are close together, but the sphere and part of the plane is between them, then the shadow will be seen. If however the planes are far apart then there will not be enough resolution in the shadow z-buffer to separate the sphere from the plane behind it and no shadow will appear. Take, for example, the case with the near clipping plane is at $n = 3$ and the far clipping plane is at $f = 78$. For the z-buffer algorithm to resolve an object the size of the sphere, we need to resolve a different z value for $z = -77$, so given

$$Z_s = \frac{f\left(1 + \frac{n}{Z_v}\right)}{f - n} = \frac{78\left(1 + \frac{3}{-77}\right)}{75} = 0.999481$$

which implies 10 bits of resolution. Therefore we will not get a shadow at all, since the z-value of the sphere is indistinct from that of the plane. Note that Z_s is a non-linear function of distance, with coarser quantization further from the viewpoint. [30%]

Assessor's comments: *(Attempted by 19/42 candidates, average mark 10.3/20). This question tested the candidates' understanding of 3D graphical rendering and the use of the shadow z-buffer. A few candidates had a surprisingly muddled view of the rôles of the modelview and projection matrices. Most of them were able to come up with a plausible value for the width of the image of the sphere, although some lost marks by calculating its area instead. Answers to the last part of the question on the operation of the shadow z-buffer were generally poor.*

Richard Prager
May 2003