

Simulation of Ground Penetration Radar

Mike Fletcher & Dick Spicer

(School of Computing, Staffordshire University, UK)

7.1 Introduction

The output from Ground Penetration Radar (GPR) surveys is notoriously difficult for archaeologists to interpret. The complexity of the signals received back from the series of radar pulses is well illustrated by Fig. 7.1, which shows the result of a real survey. Visualisation of results is limited in general to a display of each transect, either as seen here in monochrome (though frequently with fewer grey tones) or as a false-colour image, using the arbitrary colours of a default palette.

As part of a detailed study of radar visualisation, the work presented here is a computer simulation of a radar traverse, with two main aims in mind: to enable archaeologists better to comprehend radar pictures; and to provide a set of "master" images derived from simple "objects" for experimentation with image-processing methods to clarify radar images. This paper will describe the simulation and illustrate its use with a number of simple "objects".

7.2 Principles of radar survey

Before discussing the simulation, it is necessary briefly to explain the principles of the usual methods of radar survey.

Surveys are normally carried out as a sequence of traverses, spaced at some interval (typically one metre). The antenna is wheeled or dragged across the site at a more or less steady rate, sending out pulses and receiving echoes as it goes. Fig. 7.2 shows the transmitter/receiver combination at two positions (A and B) during a traverse. At position A, it is shown vertically above a single small object buried in a uniform matrix. The signal is reflected and returns vertically to point A. This produces an image (a) on the instrument output, shown on the right. The depth of this image at (a) is proportional to the time taken by the




Figure: 7.1




Figure: 7.2




Figure: 7.3




Figure: 7.4

signal to travel to position (a) from object A. Later, when the antenna unit has moved to position B, its distance from the object is greater; therefore the time taken to return to B is longer, and so the image appears at (b), which is lower because of the extra time taken.

It must be emphasised that the speed at which the antenna moves along the ground surface is very much slower than the speed of the radar signal, and also slow in comparison to the repetition rate of the pulses. The rapidity of the pulses means typically that they are spaced apart by only a couple of centimetres of travelled horizontal distance.

The complete image after the antenna has passed over the object is shown in Fig. 7.3. The shape of this curve can be shown to be a hyperbola, and its width will depend upon the angle of the beam of the radar antenna, as well as the depth of the resulting trace.

So far we have assumed that the radar signal is just a short pulse. It is, in fact, a burst of a radio wave, with a typical characteristic shown in the left side of Fig. 7.4. Hence the reflected signals will result in an image similar to that on the right of Fig. 7.4.

The picture resulting from a real radar survey is a graphical plot with the horizontal axis proportional to the number of pulses over a given time (and therefore, if we assume a constant speed, proportional to traverse distance), and the vertical axis proportional to echo delay. Colour (or intensity) represents the strength of the echo at any particular time. The image is often thought of as a vertical "slice" or section cut along the direction of the traverse, though this is a very simplistic view, as we shall see.

Leaving aside for the moment the issue of the vast imbalance between inter-traverse distance and the rate of echo-sampling along each traverse, there are many very considerable characteristics of the echo behaviour which complicate the pictorial representation which the archaeologist sees. We list below some of the more important properties which affect the image:

1. The distribution of the antenna, both during transmission and reception, is directional and conical (with lobes), with the strongest signal along the main axis (usually vertically downwards). Both the transmitted signal and the received signal reduce, therefore, as the angle from the axis increases.
2. The signal is absorbed and scattered by the material through which it passes, depending on the nature of that material.
3. Interfaces between different materials encourage reflection (and refraction) of the signal.
4. The velocity of the signal is dependent on the dielectric constant of the material through which it passes, being fastest with materials of low values (and very fast through air).
5. The pulse emitted by the transmitter has a characteristic attack and decay, and is, of course, an alternating signal.

There are other properties, such as different penetration with different wavelengths, which were not considered to be relevant to this particular modelling exercise.

7.3 Modelling ground penetration radar on a computer

7.3.1 Properties of the program

There were a number of reasons for modelling GPR on a computer: the first was to help clarify and visualise — both for us and for others — what each of the above factors actually did to a survey; the second was to demonstrate how certain features could (or could not) be perceived and hence interpreted by the survey, and whether there was the possibility of more than one cause for any one kind of output; and thirdly, we wanted to obtain experimental data which might assist in choosing "decoding" functions for image enhancement.

The properties of the program were made individually switchable, with all parameters adjustable, so that the be-

haviour of each characteristic could be seen separately. Thus it is possible to build up a sequence of pictures, from the same input data, which shows how each characteristic of the signal contributes to the final picture. Although the user of a real radar instrument might not be able to control those characteristics which are intrinsic to the physical properties of radar, switching them on and off in the simulation is a powerful aid to understanding.

The user of the program can design any collection of features, both layers and objects, by drawing them as with a paint package. The colour or grey-level determines the dielectric "density" of the material, with black being the least "dense", and white the most. No provision for calibration of dielectric, nor of absolute scaling of dimensions has been made, because it was considered rather inappropriate given all the possibilities of survey strategy; different radar frequencies are not accommodated for the same reason. We should point out also that our model is limited to two dimensions: modelling the radar pulse and objects in three dimensions was considered too costly in computer time and would create too many difficulties for potential experimenters to visualise (let alone interpret).

7.3.2 The effect of the parameters

Taking the simplest parameters first, we will illustrate the effects they have on a simple input image, then we will add "features" to show their characteristic output image.

We can begin by assuming that the signal is completely reflected once only from the interface between one material and another which is of a different dielectric constant. Fig. 7.5 shows two objects, the one on the far left being a flat plate of no great thickness, and the other object being a fairly thick block. Fig. 7.6 shows the result of simulation, assuming that the signal is infinitely narrow, with no dispersion, and that the pulse is infinitely short. We see merely a line indicating the top and bottom surfaces of our objects. In this particular example, the upper face is shown in white, and the lower in black, so as to distinguish the two.

Note, however, that the lower (black) line in the simulation of the block appears to be lower than the position of the interface in the original image. The reason is because our "object" is of a high dielectric material, and the signal travels slower through it. Since the y-axis is proportional to the time taken for the echo to return, it appears at a lower position.

7.3.3 A more realistic simulation

So far, this has not been a realistic simulation. We know that the signal is not infinitely narrow, but conical and tapering. Fig. 7.7 shows the same traverse, but with a 30 degree antenna dispersion. Note that in the output the object's width has increased, and that the increased portions droop downwards to form "cusps". The reason is that when the antenna is not directly above the object, the distance that the pulse travels is greater than the vertical distance, as described above. A horizontal surface will be slightly more thickened for the same reason.

Reflections from an infinitely short duration pulse would appear like this, as a thin line. But in reality the pulse lasts a finite time, with a sharp leading edge, and a gradual oscillat-




Figure 7.5




Figure 7.6




Figure 7.7




Figure 7.8




Figure 7.9




Figure 7.10

ing decay. This simulation permits different characteristics of decay, though this adjustment is not particularly important to our present purposes. The image of Fig. 7.8 shows a decay covering some half-dozen signal cycles. Now the result is beginning to look rather more like the confused image we expect to see from radar. Note how the signal variation is strong near the top of the echo, and becomes weaker. A surveying instrument would allow the user to vary the gain, or sensitivity, of the received signal according to time (and

hence image depth). Too high a gain, and the upper portions of the reflection would saturate (or overload). Too low a gain results in earlier decay of the signal.

The input image has so far been very simple. If we add a surface low down on the image (Fig. 7.9), the effects of shadowing can be seen (Fig. 7.10). Because part of the original signal has been reflected by the upper object, there is less signal available to be reflected off the lower surface. Note the width of this shadow, due to the 30 degree beam




Figure 7.11




Figure 7.12




Figure 7.13




Figure 7.14




Figure 7.15




Figure 7.16

width. More noticeable, however, is the diffraction effect caused by phase changes in the signal as it passes through the upper objects.

Speed of the signal is proportional to the dielectric constant of the material. So far we have assumed complete contact between the antenna and the ground. In Fig. 7.11 a ground surface has been added, with a small depression near the middle of the picture. In addition, another surface has been at the bottom right of the picture. The result of the

survey (Fig. 7.12) shows how the low dielectric of air has "pulled" the surface right up to the top of the image. At the same time, the depression on the surface is revealed as a bump in the lower surface "floor", and the cusp of the object near the middle has been upturned. Fig. 7.13 shows in addition a dark area (representing a void or chamber filled with air). The upper surface of the chamber appears as expected, because the material above it is uniform. But the lower surface does not appear at all. In fact it does, but

because the signal velocity is so much greater through air, its reflection appears almost at the same place as the upper surface. Depending on circumstances, the double reflections may have differing phase relationships, and may reinforce or even cancel each other, making interpretation extremely difficult.

It is worth noting, too, that the lowest surface in the original picture has been raised by the same phenomenon, but in a different manner than the surface of Fig. 7.12: in Fig. 7.14 there is an interruption in the shape of the lower surface caused by diffraction effects. This demonstrates quite clearly how one should look for evidence in places other than just the immediate area of interest.

Adding a vertical feature — a wall, in the example of Fig. 7.15 — shows another very intriguing effect: the top surface of the wall is easily visible, but its sides are very indistinct (Fig. 7.16). A few moments with pencil and paper will show that vertical surfaces are only “visible” when the antenna is not directly above, circumstances when the signal is weak. In fact, the reflection becomes stronger as horizontal distance increases, but the source signal becomes coincidentally weaker.

7.3.4 Limitations of the simulation

We have shown how various properties of the radar wave result in different characteristic images, and how it is possible to have inconclusive or ambiguous results. There are still some limitations of the simulation as it stands. The one

main shortcoming is that only two-dimensional information is modelled; that is, the program does not allow for one to place an object in front of or behind the traverse. There are a number of reasons why this is not desirable. The first is that the program already takes a long time to run (a full-blown simulation might run for half an hour), and adding the third dimension would increase its complexity so as to make it unusable. Second, the complexity of mentally visualising what is going on would, we believe, hamper the main reason for the program, namely to assist understanding of the output images. Thirdly, as was pointed out above, transects are normally taken at intervals far in excess of the x-resolution of the output image. Any small object at a distance of, say a metre, would not significantly alter the output.

7.4 Conclusions

We believe that this simulation is a useful tool for hypothesis testing, for education, and for the “calibration” of the “experienced eye” of interpreters. By comparison of before and after images, we might improve our skills in interpretation, and alert ourselves to those circumstances which give rise to ambiguities. The simulation might also be a prelude to the construction of a real standard calibration test site for evaluation of new instruments.

Mike Fletcher and Dick Spicer
School of Computing, Staffordshire University.