BYU SAR
A Low Cost, Compact Synthetic Aperture Radar

A Thesis

Presented to the

Department of Electrical and Computer Engineering

Brigham Young University

In Partial Fulfillment

of the Requirements for the Degree

Master of Science

by

Bryan R. Jarrett

August 1994

This thesis, by Bryan R. Jarrett is accepted in its present form by the Department of Electrical and Computer Engineering of Brigham Young University as satisfying the thesis requirement for the degree of Master of Science.

David V. Arnold, Committee Chairman

David G. Long, Committee Member

Date

Wynn C. Stirling, Graduate Coordinator

Acknowledgements

I thank my research advisor Dr. David Arnold for his tutoring, suggestions and guidance in problem solving. I thank both Dr. Arnold and Dr. David Long for their assistance in the writing aspect of this thesis and the opportunity to perform the work. Many of my fellow graduate students contributed to the successful completion of this thesis. In particular, Ryan Reed contributed valuable advice and aided with computer code. He, Scott Collyer and Doug Daum aided with the RF subsytem. Doug Daum and Larry Arnett helped considerably with the deployment of the SAR. I am grateful to NASA for the funding which made this project possible. Lastly, my wife, Eileen, helped in the writing, test deployments and encouragement.

Contents

List of Figures

Chapter 1

Introduction

Synthetic Aperture Radar (SAR) has been around for forty years and its effectiveness as a high resolution imaging radar is well accepted. However, with limited availability and high cost of SAR imaging systems, SAR has not demonstrated its full potential as a widely applicable and effective research tool. It was the goal of this research to build a SAR system, competitive in performance with current SARs, but with a cost that makes it more accessible. This goal was met by constructing an innovative and compact SAR at a hundredth of the cost of current SAR systems. This prototype SAR system is hereafter referred to as the BYU SAR.

At this time there are few civilian SAR systems. Synthetic aperture radar has been used for military applications, but these systems are unavailable for private use. There are several systems owned by large companies such as the Jet Propulsion Laboratory (JPL) and the Environmental Research Institute of Michigan (ERIM). The cost of these systems (5 to 10 million dollars) puts them out of reach for many groups, but with an innovative design and less transmit power, this cost could be dramatically reduced.

In this thesis it is shown that current off-the-shelf technology can be applied to produce a low cost yet effective SAR. Digital to analog conversion and sampling technology are used to process the basebanded signal digitally without an intermediate frequency (IF) signal. This contributes greatly to bringing the cost of the system down to approximately $70,000, without sacrificing system performance.

The performance of the BYU SAR was demonstrated by creating an image of a field in northern Provo, Utah. In this image, a resolution of 0.45 meters was obtained in the direction of motion (azimuth direction) and 3 meters in the direction perpendicular to motion (range direction). This resolution exceeds the resolution of the JPL SAR. Although it was a little less than the resolution of the ERIM SAR in the range dimension, the corresponding sidelobes are lower and in azimuth resolution BYU SAR

surpasses the ERIM SAR.

The BYU SAR has the ability to successfully map a sample site. For example, the fence posts and house show up well. In addition to these man-made objects, trees and ground relief are easy to identify. The observable ground relief is the foundation of the house which will be evidence that the house was there even after the house is gone.

Archaeology is a field that has not been able to make full use of SAR technology because of the cost of SAR systems, even though it has been demonstrated that SAR imaging can be valuable in archaeological studies. A SAR using a carrier frequency of approximately one GHz has the ability to penetrate vegetation canopy and dry sand. Adams et. al. (1981), using the JPL SAR, made use of the ability to penetrate vegetation canopy in imaging the Mayan lowlands. As a result of this deployment, the archaeologists were able to map the extensive canal system used by the Mayans and gain a closer estimate of the size of the Mayan population. Elachi (1982) produced SAR images with data from the Shuttle Imaging Radar (SIR-A) that showed dried-up river beds at the base of the Sahara desert. Stone age artifacts were discovered when these riverbeds were verified. The resolutions of these systems found useful in archaeology were only 20 and 40 meters respectively in both range and azimuth dimensions.

With regards to the Mayan expedition, Irwin Scollar (1991 pg 585-586) said that the difficulties with this type of experiment were obtaining the SAR equipment and the poor resolution of the systems. The BYU SAR has improved resolution and, with less power, a cost that makes it accessible. Its resolution is improved over these previous systems by a factor of seven in the range direction and forty-four in the azimuth direction. Thus, the BYU SAR is an answer to Irwin Schollar's criticisms.

In this thesis, basic topics defining the BYU SAR are examined. Chapter 2 discusses how SAR can be an effective tool in archaeology and describes experiments that verify its usefulness in this field of study. An overview of basic scattering theory is also included in Chapter 2. Chapter 3 shows how SAR works and describes the basic SAR algorithm. The BYU SAR system is described in Chapter 4, giving an overview of how the SAR algorithm is implemented in BYU SAR. It also describes the construction of the SAR. The results are shown in Chapter 5 by analyzing the image taken of a local area. Many interesting features are identifiable including a house foundation. The success and difficulties of the BYU SAR are elaborated on in the final chapter.

The contributions of this thesis are 1) The hardware was built at the system level for a SAR that is 1/100th the cost of existing systems, 2) Near real-time software was developed for controlling the system, collecting and processing data, 3) The system performance was demonstrated.

Chapter 2

SAR Application

A SAR system operates at microwave frequencies, enabling it to extract different types of information about a target than is available from the visible spectrum. In order to interpret the SAR image, it is important to understand the interaction of electromagnetic waves with various materials, particularly at boundaries. Electromagnetic waves in the microwave spectrum are reflected back from interface boundaries according to a material's roughness as well as its electrical properties. Thus, both the texture and electrical properties of a target are important. Radar waves are not completely reflected by vegetation canopy and some of the waves can penetrate all the way to the ground and be reflected back. Radar is also able to penetrate several meters into some materials, like sand, that are opaque in the visible spectrum of EM waves. These last two properties make it possible for the radar to image objects through vegetation and ground cover and make SAR a valuable tool in archeological research.

Synthetic Aperture Radar (SAR) has been successful in remote sensing archaeological sites (see Adams et al. (1981), Elachi (1984) and Holcomb (1992)). Until recently, archaeological uses of SAR have been limited, primarily due to the cost of the SAR systems and their limited resolution. The BYU SAR is an experimental radar system that provides good resolution at low cost and can be a useful tool in archaeological research.

2.1 Archaeological Applications

When archaeologists have a site well mapped out, they have a better perspective on the size and diversity of a civilization. Site mapping is frequently done with aerial photography, but in some locations a perennial cloud cover or a vegetation canopy prohibit this type of mapping. Site mapping can be done from the ground, however, thick vegetation can cover landmarks and make it difficult to move surveying and mapping equipment. Radar overcomes both of these obstacles and can accurately

record site information. Since microwave radiation below 20 GHz penetrates cloud cover with little attenuation, observation can be done regardless of the weather. A radar operating below about 4 GHz can penetrate vegetation canopy, making it possible to image objects on the ground below the canopy.

SAR has already been used on archaeological sites, but only on a limited basis because of the cost of the systems involved. It has successfully been used to penetrate vegetation canopy and sand to image objects of importance to archaeologists with an inherent mapping scale. The two examples are in the Mayan lowlands and the Sahara desert.

Adams et. al. (1981) used SAR in Central America on the Mayan lowlands to penetrate the thick vegetation. The archaeologists wanted to know the extent of the irrigation system used in the swamp areas in order to understand how the Mayans provided food for their population. This was difficult to discern because the thick vegetation prohibited aerial photography and conventional ground mapping techniques. The JPL SAR, with a carrier wavelength of 25 cm (1.2 GHz) and a resolution of 20 meters, successfully identified ancient canals and fields. The SAR provided new information about the economic base and agricultural practices of the ancient Mayan population. Better resolution could improve its usefulness since a broader range of structures will be detectable and identifiable.

Elachi (1984) was successful in imaging the base of the Sahara dessert through an average of 2 meters of sand using the SIR-A SAR. The image showed the existence of large river beds. When the river beds were excavated and verified, numerous Stone Age artifacts were discovered (Holcomb). The wavelength used was 24 cm (1.25 GHz) with a resolution cell size of 40 meters.

Large man-made objects buried underground tend to be bright. Their corners reflect much of the radar signal, and many man-made objects contain metal which has a bright return, contributing to their detection. These attributes help them show up well in a SAR image.

As mentioned earlier, the two biggest hindrances for the use of a SAR system in archeology have been cost and resolution. Irwin Scollar (1988) comments on the use of SAR in the Maya lowlands study: "This experiment is unfortunately unique, for on the one hand it is not easy to obtain equipment of the type available to the Jet Propulsion Laboratory and on the other hand there are few problems in archaeology which require the detection of fossil relief with a resolution of the order of tens of metres." Current SAR systems typically cost millions of dollars. A SAR system costing under $100,000 would make SAR much more accessible.

Instead of building a SAR, it is cheaper to use data that is available from satellite SAR deployment. However, this is a hit and miss solution that does not cater to a particular project's needs. The resolution of the images may not be appropriate for the desired use. The desired site might only be

partially imaged, or not imaged at all. The time when an image is made may also limit its usefulness.

Resolution of around twenty meters can be useful in locating large objects like the river beds in the Sahara. Some of these large objects are associated with population centers and artifacts can be found in the area. This is the case in the Sahara. With better resolution, it is possible to identify buildings since the corners would be characteristically bright. Thus, there is a need for an affordable SAR system that has a resolution of around 2 meters.

The BYU SAR system can fill these needs in the archeological field, though for this application, an L-band system would be more effective (BYU SAR was built as an X-band system). It is a low cost prototype system with good resolution and enough transmit power and receive gain to be deployed from a small low-flying aircraft. The hardware cost for this experimental system was about $70,000, but implementing a system for archaeology will be less expensive since RF parts would be cheaper for an L-band system. The present BYU SAR system is designed for a resolution of about 2.5 meters, which is much better than most of the systems implemented thus far. The wavelength can be chosen to penetrate either vegetation or dry soil to identify structures.

2.2 Scattering Theory

Interpreting SAR images requires a knowledge of scattering theory. Scattering theory describes how EM waves interact with the boundaries of different materials and the materials themselves. This knowledge brings an understanding of what makes objects appear bright in a SAR image, and whether a bright portion of a SAR image is from the surface of the target, or possibly buried. If it is buried, we will have an idea of how deep the bright object could be. The use of scattering theory enables interpretation of the SAR data.

When an electromagnetic wave is incident upon a surface, some of the energy is transmitted into the material and some is reflected back. The reflected wave is said to be scattered. The wave transmitted into the material may be scattered later at another interface or from inhomogeneities in the material. These two types of scattering are called respectively surface scattering and volume scattering.

The first type, surface scattering, can be broken down into two nonexclusive types, specular and diffusive. Specular scattering occurs with roughness scales that are much larger than the wavelength. It is similar to visible light reflecting off a shiny surface; it only reflects back at you if the normal to the surface is pointing toward you. The second type of surface scattering, diffusive scattering, occurs when the size of the surface roughness is closer to the size of the wavelength. In this case, the scattering occurs

in many directions while only the portion being scattered in the specular direction is detected. Scattering from a target can be a combination of both specular and diffusive, caused by roughness of different scales on the target (Ulaby et al chapter 11).

The amount of electromagnetic energy that is reflected depends on the material permittivity and conductivity. Reflection and transmission at the surface of a material with very low conductivity are determined by the difference in permittivity of the two materials. A large difference in permittivity causes high reflection while a small difference results in high transmission (and therefore, small reflection). An example of this is imaging the bottom of the Sahara desert. The relative permittivity of air is about 1 and the relative permittivity of sand is about 2.5, resulting in a small reflection. The rock bottom beneath the sand has a relative permittivity of 8 which causes a high reflection. This enables the radar to image the desert bottom.

Conductivity is also important. A perfect conductor reflects all of the energy and transmits none. As the conductivity of the target goes up, its coefficient of transmission goes down and more power is reflected. An example of this is damp soil. As the soil moisture content goes up, its conductivity increases and it reflects a greater percentage of the electromagnetic energy.

In addition to the percentage of the electric field transmitted, the amount of attenuation the waveform will experience in a homogeneous medium is important. Attenuation (or loss) is experienced because of the conductivity of the medium in which the wave is travelling. The radar energy is dissipated in the form of heat. A common way of describing the losses in a medium is to define the distance traveled before the electric field is reduced by the factor 1/e or, to about 37% of its original value. This is called the penetration depth, δ. It is proportional to the square root of the wavelength of the carrier used and is dependent on material properties. The higher the conductivity, the shorter the penetration depth, and, therefore, the higher the attenuation. The same relationship can be seen with the carrier frequency; that is, higher frequencies experience higher attenuation. This is why successful subsurface imaging in dry sand requires lower frequencies.

An example of subsurface imaging is the work done by Elachi. The penetration depth in dry sand with a wavelength of 30 cm (1 GHz) is 2 - 8 meters. Some penetration can be achieved at low frequencies in slightly moist soils since the conductivity of the soil does not increase as drastically until the water content rises above 15% in volume. With a water content of 15%, a similar penetration, 1.5 meters, can be obtained if a 300 meter wavelength is used (the radar frequency must be 1 MHz) (Balanis pp 208-209).

The refraction of radar waves in sand helps subsurface imaging. The EM waves that are

incident on the sand surface are bent to a steeper angle because of the higher index of refraction of the sand. This steeper angle provides a shorter in-sand path length and therefore the attenuation is less than it otherwise would be.

The second type of scattering is volume scattering (Curlander pp 55-57). Electromagnetic waves in an inhomogeneous media experience volume scattering. This scattering comes from pockets of material with differing permittivity or conductivity. An example is a vegetation canopy as seen in the Mayan lowlands. In the vegetation canopy, the media is air with pockets of vegetation. The permittivity and conductivity of the vegetation are much higher than that of the air, so each leaf and branch is a scattering surface. The scattering that takes place in this case is dependant on the difference in permittivity and conductivity, the size of the scatterers relative to the wavelength, and their distribution and orientation. In general, the EM waves with longer wavelengths have better penetration.

The Mayan lowlands experiment is an example of both surface and volume scattering. The radar wave traveled through the air with little attenuation until it reached the vegetation canopy where it experienced some surface and volume scattering. Part of the wave was scattered back while most penetrated to the ground, where surface scattering occurred. Most of the EM wave incident upon the ground was reflected because of the high conductivity of the moist soil. The reflected wave was again scattered, and therefore attenuated, while traveling through the canopy the second time before reaching the receive antenna.

Surface scattering occurs at the boundaries of two different materials. If we know a material's conductivity and permittivity, we can anticipate whether or not the radar will penetrate to a significant depth and possibly be scattered by an object below the surface. With this knowledge, meaningful information can be retrieved from radar data.

A common problem for SAR images is speckle noise. Speckle noise is caused by constructive and destructive interference of coherent electromagnetic waves. This is not seen in optical images because of the incoherent light source. The waveform scattered by a given target inside a resolution cell is added with the waveforms generated by scattering from all the other targets and surfaces that are in the same resolution cell. The waveforms add constructively and destructively, depending on the antenna position and the irregularities in the targets close to the same size as the wavelength. Because of the interference of the localized points of the terrain the measured backscatter is considered to be a random variable. This randomness of the measured backscatter is called speckle noise and is a characteristic of SAR images. Speckle noise causes homogeneous scenes to have a fuzzy or speckled appearance. This must be taken into consideration when analyzing a SAR image.

Chapter 3

Basic SAR Overview

Synthetic Aperture Radar is an imaging radar operated from a moving platform such as a satellite, an airplane or a moving vehicle (Sherwin et al). The imaging geometry is shown in Figure 1. If the radar transmits pulses of a single frequency, then the resolution in the direction perpendicular to motion is determined by the time width of the pulse transmitted, or the instantaneous ground illumination, as shown in Figure 2.

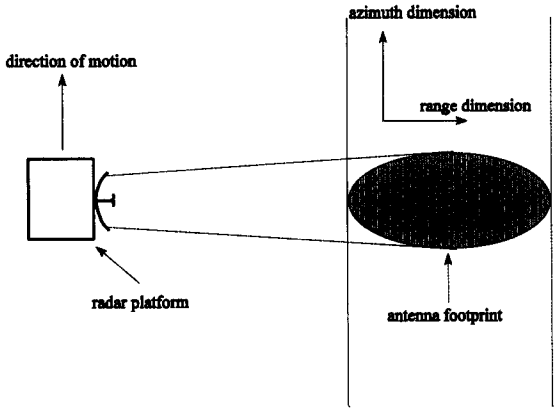

Figure 1: SAR deployment diagram.

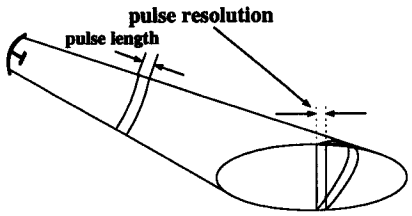

Figure 2: Resolution in the range direction.

Resolution in the direction of motion or azimuth direction is achieved by Doppler filtering of the return. The Doppler frequency shift is due to the relative motion between the radar and a target. A target ahead of the radar has a positive Doppler frequency. This frequency decreases and reaches zero

when the target's position vector is perpendicular to the radar's velocity vector. The Doppler frequency becomes increasingly negative as the target gets farther behind the radar. The azimuth resolution of a target is achieved by filtering out all the Doppler frequencies that do not correspond to the target's location for each target. In a typical SAR system, this resolution can have a cell size on the order of one meter. This high resolution imaging is useful in remote sensing of the earth and has been used to image ocean waves for ocean wave spectrography, polar ice formation and basic land mapping (Bruning, Drinkwater).

3.1 Pulse Compression

High resolution is achieved in SAR in both the azimuth and range directions, as shown in Figure 1. The signal processing technique that is used to do this utilizes information we have about the signal. It compresses a long pulse into a narrow sinc function and is called pulse compression (Klauder, Farnett).

Pulse compression is based on the fact that sine waves are correlated only if they have the same frequency and phase. To prevent the correlation from being periodic (as with a constant frequency), a sine wave of continuously changing frequency is used in pulse compression so that it correlates with itself at only one time shift. The most common waveform of this type is a finite duration waveform that has linear frequency modulation (LFM), referred to as a chirp. The equation for an LFM chirp is:

$$x(t)=\cos\left(2\pi\left(\frac{k}{2}t^2+f_0t\right)\right)\left[u(t)-u(t-T_p)\right] \qquad (1)$$

where f_0 is the starting frequency of the chirp and k is the chirp rate. The time-bandwidth product (tB) is a useful metric in analyzing the properties of the chirp. The time-bandwidth product (the pulse duration multiplied by the pulse bandwidth) is given by $tB=kT_p^2$ where T_p is the pulse length. A typical chirp is shown in Figure 3. In this case f_0 is 0 Hz.

The autocorrelation function, Y(t), of the chirp is used in processing the pulse. It is given by

$$Y(t)=\int_{-\infty}^{\infty}x(\tau)x^*(t+\tau)d\tau. \qquad (2)$$

The correlation of the chirp is peaked at a point so the resolution is given by the width of the peak rather than the width of the chirp pulse. Hence, pulse compression results in improved resolution (refer to Figures 3 and 5).

The number of computations required to compute the correlation can be reduced by using Fast

Fourier Transforms (FFTs) and by taking advantage of the convolution and correlation properties of the Fourier transform. If $X(\omega)$ is the Fourier transform of x(t), then the Fourier transform of the correlation of x(t) with itself is the simple product

$$Y(\omega)=X(\omega)X^{*}(\omega). \qquad (3)$$

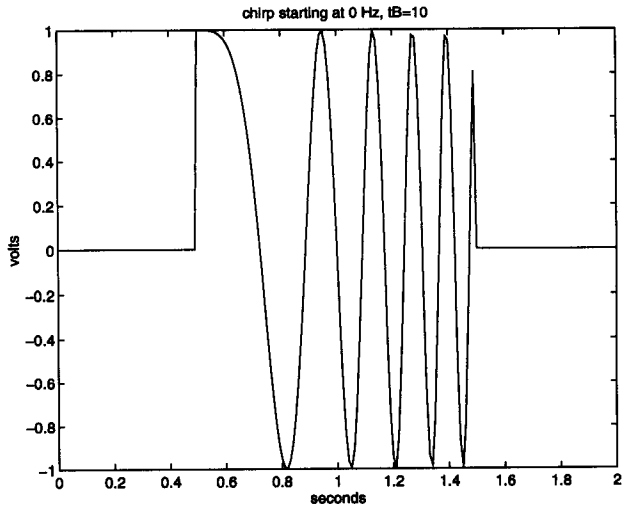

Figure 3: A chirp waveform.

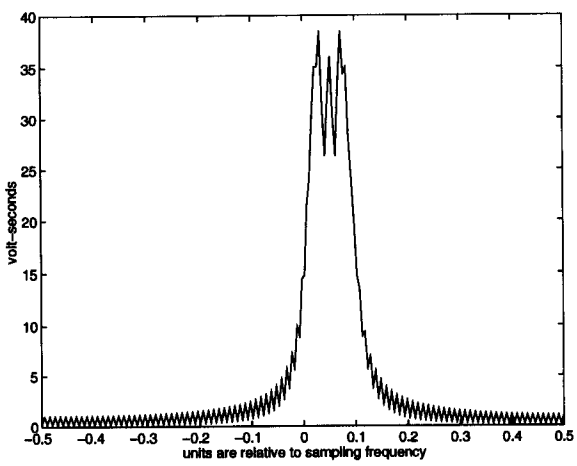

Figure 4: The spectral magnitude of a chirp with tB=2.5.

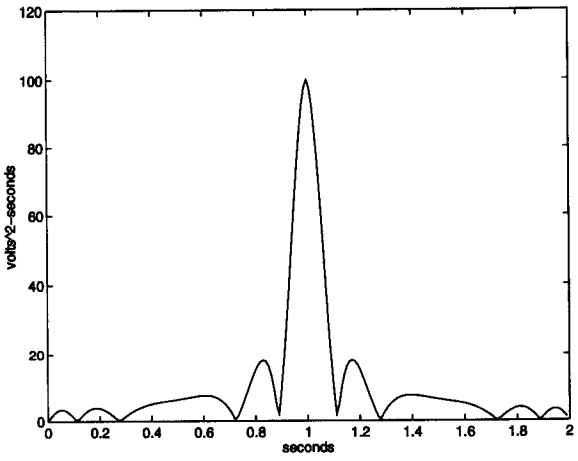

Figure 5: Autocorrelation of chirp with tB=2.5.

If the time-bandwidth product is greater than 10, the pulse width is approximately the inverse of the bandwidth. The magnitude of a chirp's Fourier transform is similar to a rect function if the chirp has a time-bandwidth product greater than about 10. This is demonstrated in Figure 6. For tB<10, this simple rule doesn't apply. See Figure 4. In Figure 5 this product was 2.5 and the correlation function is wider than 1/B. Figure 6 shows the magnitude of the Fourier transform of a chirp with the same

bandwidth as shown in Figures 4 and 5, but with a time-bandwidth product of 100. Figure 7 shows its autocorrelation. The autocorrelation of the chirp with a large bandwidth and sufficient time-bandwidth product is much closer to the ideal of a delta function, but there are sidelobes that do not extend past the first sidelobe because T_p is so short that there is minimal overlap in the correlation of the low tB chirp. These sidelobes are undesirable.

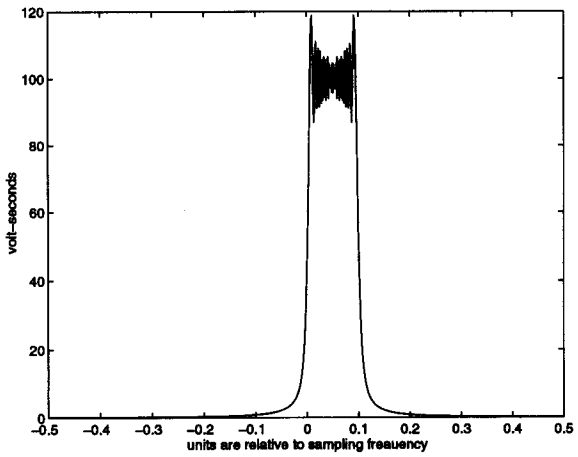

Figure 6: Spectral magnitude of chirp with tB = 100.

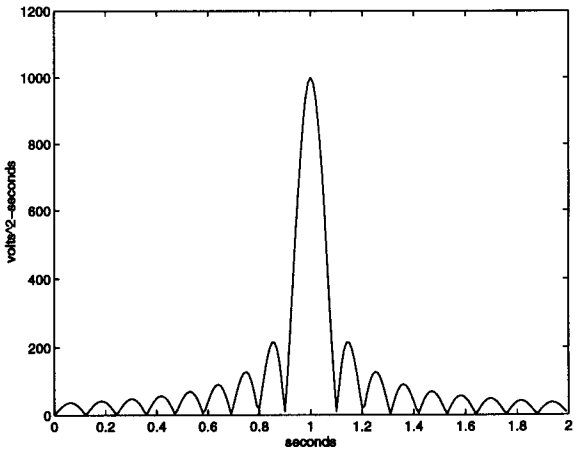

Figure 7: Autocorrelation of chirp with tB = 100.

Effective use of windowing on the transmitted pulse can increase the quality of an image by minimizing sidelobes. The sidelobes from one target interfere with the mainlobes of other targets. This interferance can be either constructive or destructive and can change the amplitude or even split the response of a target. Having a split in the mainlobe of a response causes an artifact since it appears to be two targets instead of one. These effects can be minimized with windowing (Harris), but at the expense of resolution. An effective window's sidelobes will be below the quantization noise level of an image. This increases the 3 dB mainlobe width of the correlation function by about 50%. Windowing is easily accomplished in the range direction because the two chirps that are correlated are under the control of the transmitter and therefore can be designed with any desired windowing. In azimuth correlation, the antenna pattern provides the window.

3.2 The Matched Filter

An alternative view of pulse compression is based on a matched filter interpretation. A signal, x(t), that is filtered with a filter having an impulse response, h(t), can be represented by the convolution of x(t) with h(t),

$$x(t)*h(t)=\int_{-\infty}^{\infty}x(\tau)h(t-\tau)d\tau. \quad (4)$$

The matched filter is the ideal processing technique for a maximum signal to noise ratio in the presence of additive Gaussian noise. It can be shown that the optimal filter response in this case is (see the appendix for the derivation)

$$h(t)=x^{*}(-t). \quad (5)$$

This can be interpreted as the autocorrelation of x(t). The obvious shortcoming of this technique is that the signal must be known before the filtering can be done, but in the case of radar this is not a problem since the return echo is a time and frequency shifted copy of the transmitted signal.

A key issue in pulse compression is resolution. The resolution can be improved by increasing the bandwidth of the broadcasted chirp, s(t), as can be seen by examining the response of the system. The time domain response, r(t), of the system is simply the noiseless matched filter response (Curlander pg 132)

$$r(t)=\int_{-\infty}^{\infty}|S(f)|^{2}e^{j2\pi ft}df. \quad (6)$$

If we model the spectral content of the broadcasted signal as a constant over a fixed bandwidth with arbitrary phase, the integral becomes the inverse Fourier transform of a rect function. As mentioned, the degree to which the magnitude of a chirp's spectrum approximates a rect function is dependant on its time-bandwidth product. The time domain response of a rect function with bandwidth B is a sinc function having a zero crossing width of $\delta t \approx 2/B$. This is the resolution in the time domain between the nulls of the sinc function, but in practice the 3 dB width is used which is roughly 1/B.

3.3 Range Compression

Range compression is used to obtain resolution in the direction perpendicular to the velocity vector. This may be done by transmitting a very short pulse, resulting in a resolution size proportional to the pulse's physical extent in space (see Figure 2). The difficulty inherent in this method is that it requires a high power pulse. Such pulses require expensive transmitters. However, pulse compression can produce similar resolution without a high power transmitter. The same energy is transmitted, but is spread over time so the peak power is reduced.

Using range compression and the round trip time of a transmitted waveform, targets are mapped in the range dimension. When the waveform x(t) is transmitted, it is reflected by a target with a brightness distribution d(t). The return waveform is the correlation of x(t) with d(t) plus a noise

component n(t). In this case, d(t) is the total brightness at a given range in the antenna footprint specified by the time it takes the signal to travel the round trip to the target. Thus $t = 2x/c$, where x is the distance to the antenna and c is the speed of light. A point target will give a return of $\alpha x(t-\tau)$ where α is a scaling factor related to the target's brightness and τ is its position in time. If we transmit a chirp and use pulse compression, we will get a sinc function centered at τ having a width inversely proportional to the bandwidth of the chirp. The noise is white and uncorrelated with the signal so its effect on the correlated waveform is minimized by the correlation integral (which is the matched filter). Correlation is a linear operation so the return from a distributed target will be the convolution of a sinc function with the terrain distribution. As the sinc function gets narrower, it approaches a delta function and a delta function convolved with the terrain brightness gives the terrain brightness.

3.4 Azimuth Compression

As previously discussed, azimuth resolution is obtained by Doppler filtering. This can be viewed as a form of "pulse" compression, except in the azimuth direction. Azimuth compression is very similar to range compression. The big difference is that the waveform used in the matched filter for azimuth compression is a result of Doppler frequency shift instead of a transmitted chirp and is dependent on a number of system parameters not in direct control of the designer. The Doppler frequency of terrain for a moving target with speed V is dependent upon the angle, θ, from the axis perpendicular to the direction of motion and is given by

$$f_d = \frac{2V\sin(\theta)}{\lambda} \approx \frac{2Vx}{\lambda R} \tag{7}$$

where R is the range and x is the azimuth component of the distance from the point to the antenna. The approximation is good when θ is small and the actual distance to the target is close to the range component of the distance. Thus the Doppler frequency varies linearly with respect to x and so can be modeled by correlation with a chirp. From our previous results we see that the resolution will be dependent on the bandwidth of the azimuth chirp. The bandwidth of the chirp will be dependent on the beamwidth of the antenna (the change in θ) and the speed, V, at which the platform moves. It is given by

$$B = \frac{2V2\sin(\frac{\beta}{2})}{\lambda} \approx \frac{2V\beta}{\lambda}. \tag{8}$$

In this equation, β is the beamwidth of the antenna which is given by λ/d_a, where d_a is the width of the

antenna. Therefore, the equation for the bandwidth becomes

$$B=\frac{2V}{d_a}. \tag{9}$$

We stated earlier that the time resolution is 1/B, then the distance resolution is the velocity multiplied by the time resolution. The azimuth resolution in distance therefore is

$$\delta_a \approx \frac{Vd_a}{2V}=\frac{d_a}{2}. \tag{10}$$

This is the optimum result and can be approached only if the time-bandwidth product is sufficiently large; otherwise the resolution is somewhat less than this. However, it is interesting to note that the resolution is inversely proportional to the antenna size. The limiting factors in this are the maximum bandwidth allowed by the sampling rate, the required power density and processing data length.

Azimuth compression is executed independent of the range compression, but it usually is performed after the range compression. This is because the data is gathered in streams ready for range processing, but many return waveforms are required to do the azimuth compression. It may also be desirable to do the range processing in real-time. Azimuth compression is performed on points of constant range. The length of the data is set is equal to the number of points of the Fast Fourier Transform (fft) used, 1024 for the BYU SAR. The matched filter for azimuth compression is given by

$$h(t)=\begin{cases} e^{f_{Dc}t+\frac{1}{2}f_R t^2} & for \quad |t|<\frac{\lambda R}{2d_{aV}}, \\ 0 & otherwise\ . \end{cases} \tag{11}$$

The f_{Dc} term is the Doppler centroid. It is nonzero when the main lobe of the antenna being pointed at an angle other than 90° from the direction of motion. The Doppler centroid is difficult to determine from the setup and must be obtained from the data. The term f_R is the frequency rate given by the bandwidth divided by the length of time a target is in the antenna footprint, and in terms of system parameters is

$$f_R=-\frac{4\pi V^2}{\lambda R}. \tag{12}$$

Range and azimuth compression have different computational considerations (see Figure 8). The matched filter's length and chirp rate, f_R, in azimuth compression are both dependent on the range and therefore the filter must be calculated for each range. Recalculating the matched filter and taking its Fourier transform slow down the azimuth compression. In addition, memory access is slower for azimuth compression. Range compression can be done with a simple block transfer of data, whereas in

azimuth compression, each data sample must be located before it is accessed.

This chapter has described the basic SAR correlating algorithm or the "rectangular algorithm", (Curlander) describing pulse compression in the azimuth and range dimensions. There are other algorithms such as "deramping" and "unfocused SAR" (Ulaby chapter 8) that are useful under different circumstances. For information on these or more information about the rectangular algorithm see an appropriate text book such as Curlander (1991). In the next chapter, we shall now see how the rectangular algorithm has been implemented in BYU SAR.

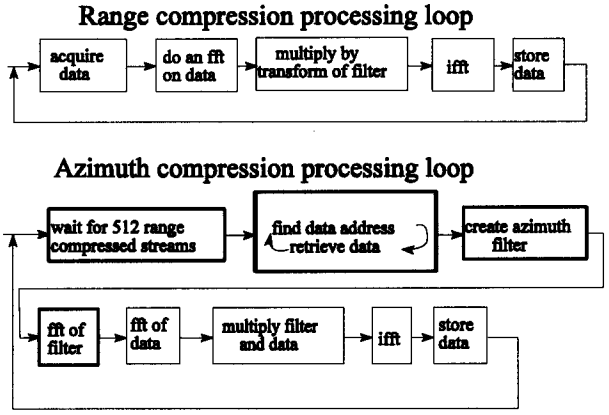

Figure 8: Flowchart showing azimuth and range processing loops.

Chapter 4

The BYU SAR System

The BYU SAR is a unique system in that it achieves good resolution at a much lower cost than other privately owned SAR systems. The system can be separated into four main elements as shown in Figure 9. First is the RF subsystem that processes the signal as an analog microwave signal. Next is conversion from a digital signal to a baseband analog signal and from an analog signal to a digital signal, in the Baseband subsystem. In this system, there isn't an IF or intermediate frequency because of on-board processing. This is the key to the BYU SAR system and makes it possible to achieve a low cost high resolution system. The antennas make up the third subsystem. Finally the signal processing is performed digitally with a personal computer in the Processing subsytem. The system as a whole is compact and relatively inexpensive. It is a competitive alternative to other SAR systems.

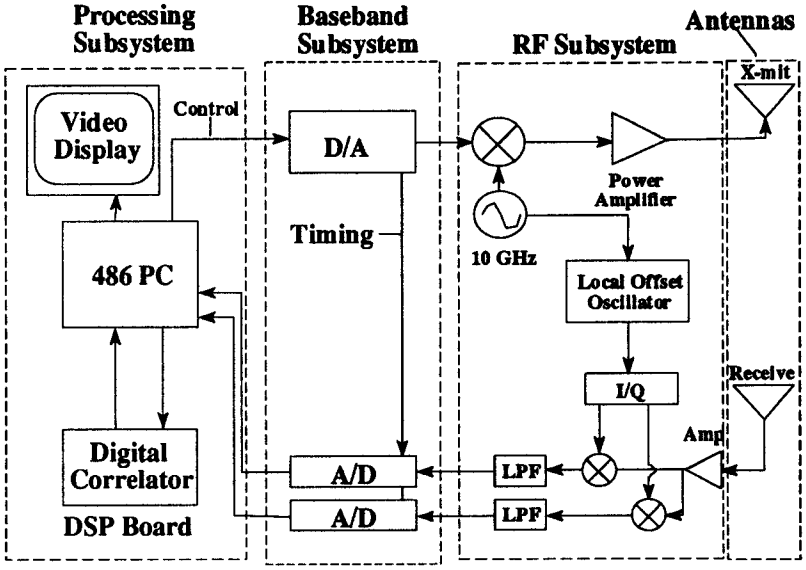

Figure 9: The SAR system.

4.1 The RF Subsystem

The first major part of BYU SAR is the RF subsystem. The RF subsystem has three basic parts as shown in Figure 9: the transmit section, the receive section and the offset local oscillator generation. The transmitter prepares the waveform produced by the D/A system for transmission through an antenna. This is accomplished by mixing the waveform from baseband up to RF and amplifying the signal. The receiver and offset local oscillator work together in mixing the RF radar echo from the antenna to an offset baseband and the receive line amplifies it so that it can be sampled by the A/D subsystem.

4.1.1 The Transmitter

The transmitter takes the chirp waveform and mixes it up to an RF frequency. It is then amplified by 30 dB for tranmission. The output power is about one watt, so the system is relatively low power. The transmitted signal is double sideband because both the positive and negative chirp frequencies are present at the RF carrier frequency as shown in Figure 10 (b). In traditional SAR systems a single sideband signal is transmitted. However, this requires more expensive hardware.

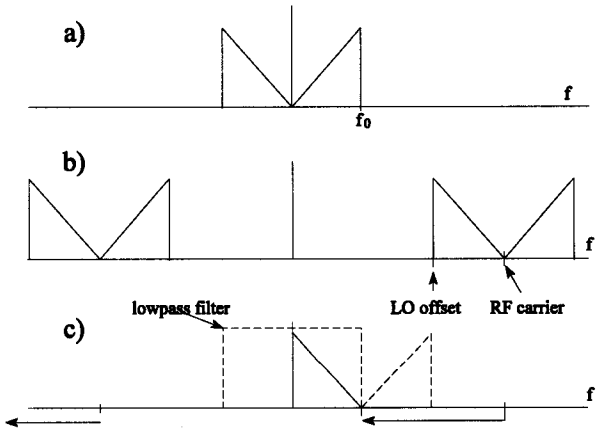

Figure 10: Diagram of (a) the chirp spectra (b) transmit spectra (c) receive mixing .

4.1.2 The Receiver

The receive section amplifies the signal by 60 dB and mixes it down next to baseband with the complex sinusoid from the offset local oscillator. The waveform is mixed down next to baseband to avoid aliasing with I/Q sampling at 100 MHz. The complex chirp has an information bandwidth from

0 to 90 MHz, but because the signal is real, its actual bandwidth is -90 to +90 MHz or 180 MHz. The return signal would need to be sampled at more than 180 MHz if the signal was mixed down to baseband. To avoid this, the signal is mixed down to an offset baseband with all the information in positive frequencies. The lower half of the chirp forms an inverted complex chirp that starts at the high frequency and decreases to zero. The other half of the chirp is then filtered out of the waveform as shown in Figure 10 (c) to produce a single sideband waveform. Because it is sampled I and Q, the signal bandwidth is 90 MHz instead of 180 MHz as shown in Figure 10 (c). Half of the power is lost in doing this, but all of the information is still present.

4.1.3 The Offset Local Oscillator

The signal is mixed down from the RF frequency in Figure 10 (b) to baseband by a multiplication with a complex sinusoid generated by the local offset oscillator. The frequency of the complex sinusoid is the RF offset by 90 MHz. 90 MHz is the maximum frequency of the transmitted chirp [f_0 in figure 10 (a)] so that the complex sinusoid shifts the RF echo down to baseband as previously decribed. 90 MHz is less than one percent of the RF frequency and it must be constant enough in frequency to allow measurement of the Doppler frequency shifts on the order of 1 Hz. This requires a high degree of accuracy in the offset LO generation relative to the RF signal. It would be much more expensive to use two sources in the RF range that are locked in their frequency difference, therefore it is accomplished by modifying the RF carrier with a 90 MHz cosine. The accuracy is then dependant on the stability of the 90 MHz cosine generation.

The offset local oscillator used in the complex bandshift is produced by using I/Q splitters and a 90 MHz signal. Both the LO and the RF signals are split to produce sine(RF), cosine(RF), sine(LO) and cosine(LO) signals. These signals are then multiplied and added as shown in Figure 11. The result is a sinusoid at the desired frequency that goes through a 90 degree phase splitter to produce a complex waveform. The power of unwanted frequencies are at least 30 dB down from the desired frequency's power so that a clean tone is produced.

4.2 The Digital and Analog Conversions

The Digital to Analog (D/A) and Analog to Digital (A/D) conversion is BYU SAR's second major subsystem. The D/A is performed by a Lecroy Arbitrary Function Generator (Lecroy AFG)

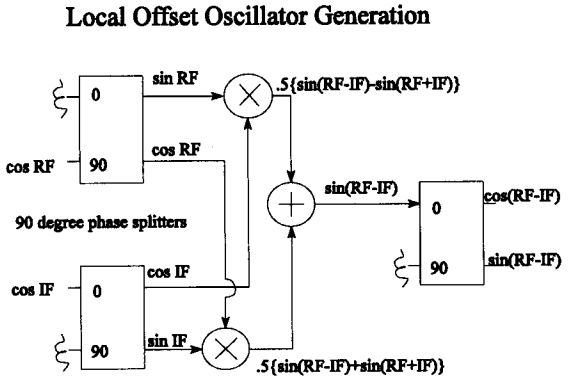

Figure 11: Offset LO generator.

which creates the transmitted chirp used in range compression. The A/D is accomplished by a set of slaved Gagescope CS250 boards that digitize the return waveform.

High performance A/D and D/A conversion at a relatively low cost enables this SAR system to be affordable and still have high resolution. The performance can be broken down into two divisions, sampling rate and the Pulse Repetition Frequency (PRF). The maximum frequency for conversions between the discrete and analog domains is the key to high range resolution and is 200 MHz for BYU SAR. The PRF, or the number of pulses that can be created and sampled per second can be a limiting factor in the azimuth resolution and is 330 Hz for this system. The A/D and D/A systems that are used provided the highest performance commercially available at the time and represent nearly half of the cost of the project. They were much less expensive than the systems that have been employed on other SARs.

4.2.1 The Digital to Analog Conversion

The Lecroy arbitrary function generator creates the transmitted chirp. A chirp is digitally created on the PC and stored in the Lecroy memory. The Lecroy AFG converts the sample points into voltages at 5 nsec intervals so it has a sampling bandwidth of 200 MHz. According to the Nyquist theorem the highest frequency that a waveform sampled at 200 MHz can have without aliasing is 100 MHz. Then, since digitally creating a waveform is equivalent to sampling a continuous time waveform, the maximum frequency of the transmitted chirp is 100 MHz. The actual transmitted chirp had a maximum frequency of 90 MHz because the Nyquist theorem requires the use of an ideal lowpass filter with gain to be able to recover the original waveform. Figure 12 shows an interpolated plot of the transmitted chirp. Because of the sampling, not all of the higher frequency phase shifts are exactly represented, but since the same waveform is used in the correlation, this is not a problem.

The Lecroy AFG is controlled over an RS-232 interface bus and provides the subsystem timing. The chirp data file is loaded into the Lecroy AFG memory over the RS-232 bus. In operating the Lecroy AFG, the data file is loaded into the function generator's fast memory and the function generator is activated by commands over the interface bus. The Lecroy AFG sends a trigger pulse to the sampling cards every time the waveform is transmitted to allow the cards to start sampling the return echo. The transmitted waveform is a series of chirps as shown in Figure 12 at a Pulse Repetition Frequency (PRF) of 156,000. A Hamming window is used on the chirp.

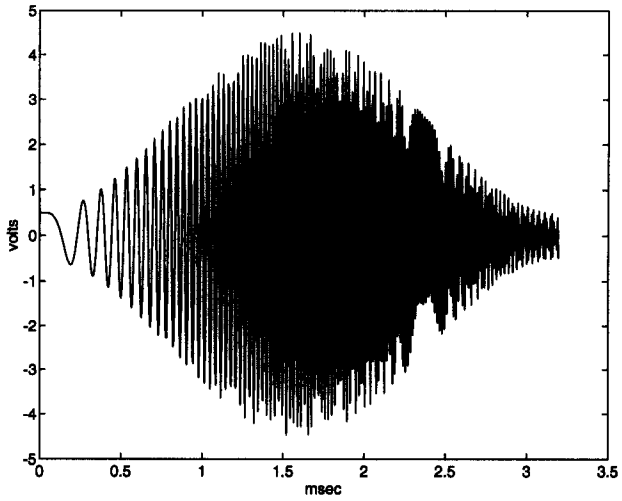

Figure 12: The transmitted chirp.

4.2.2 The Analog to Digital Conversion

The A/D part of the subsystem converts the continuous time waveform in the RF subsystem to a discrete time signal that can be processed by the digital computer. The two slaved Gagescope boards perform this by sampling the I and Q channels at 100 MHz each to provide complex samples. The cards are slaved to provide simultaneous sampling.

The sampling cards are installed in a 486 Personal Computer (PC). A combination of custom and manufacturer's code controls the cards. When the cards are initialized, they sample data at 100 MHz with 8 bit samples. When the command to get data comes from the PC, the boards wait for a trigger pulse from the Lecroy AFG. As the trigger pulse is received, the memory addresses that are being written to at that time are recorded for each card. These are called the trigger addresses. After the desired amount of data is recorded, the sampling is stopped and the data is transferred from the memory on the boards to the PC memory.

These sampling boards had the highest performance available at the start of this project (1993). It took some effort to get them to operate in a way that is useful to the SAR system. The two main obstacles were (1) getting the cards to sample a large number of pulses fast enough and (2) the uncertainty in the trigger address. The problem with sampling the pulses fast enough is attributed to the minimum time between triggering the boards. This includes time to regulate the sampling boards and transfer data. The second problem is that the last two bits of the trigger address are discarded when the boards are operating at 100 MHz so the relative trigger address between the two boards can be off by as much as three data points. When the trigger addresses are not synchronized for the two boards, the data no longer has the proper quadrature sampled phase relationship.

The minimum time between trigger addresses determines the Pulse Repetition Frequency (PRF) of the system. The PRF provides the azimuth waveform sampling. The PRF must be at least as high as the peak frequency of the azimuth chirp to allow the I and Q samples to accurately represent the waveform. The peak frequency (due to Doppler shift) is sensitive to the antenna pointing direction, antenna beamwidth and the platform speed. Because it is dependant on the platform speed, the PRF limits the speed of the radar and therefore its applications. The only way to increase the PRF is to streamline the code for the sampling boards gathering data without real-time calculations. If the range compression is not done real-time, the data processing can take place immediately after the data collection is finished and in a few minutes, an image can be examined. The PRF for the BYU SAR without real-time processing is 330 Hz which allows a maximum platform speed of 375 mph. This limits application on some airborne platforms and requires a great deal of memory. For real-time range compression, the BYU SAR has a PRF of 31.8 Hz allowing a top deployment speed of 40 mph.

4.3 The Antennas

Two X-band parabolic dish antennas are used in the BYU SAR. Traditionally, a single antenna is used, but a bistatic system is less expensive and allows a continuous wave transmission instead of the antenna switching between transmitt and receive. The transmitted chirp is broadcasted at a very high PRF to minimize the sampling board's wait for the free-running AFG's trigger. This bistatic system allows imaging of relatively close targets with an arbitrarily long transmitted chirp since sampling of the return waveform does not have to wait until after transmission is completed. To minimize the feed-through of the transmitted chirp, an attenuating material is placed between the two dishes. This allows more power to be transmitted before waveform clipping in the receiver amplifier is an issue.

4.4 Data Processing and the Final Image

The data processing is the final subsystem of BYU SAR. Its purpose is to convert the raw data into an image with useful information. The SAR correlating algorithm discussed in Chapter 3 is implemented on a DSP board and a final greyscale image is produced on the PC monitor.

The processing is structured to allow real-time implementation of range compression and display. To do this requires fast processing of floating point values, so the TMS320C30 digital signal processing board was used. Without near real-time impementation requirements, the TMS320C30 dsp board is unnecessary. Range and azimuth compression are both done on the DSP board. Range processing is done first and a magnitude image of the range compressed data is displayed as the data are processed. Azimuth compression is performed on the data after a sufficient number of waveforms are range compressed. The data are then converted to a magnitude format and the fully compressed image is displayed on the PC's monitor.

The architecture of the TMS320C30 digital signal processing board is well suited for SAR data processing. The DSP board performs a floating point multiplication in a single cycle. It is designed for doing FFTs, the biggest time consumption in range processing. Using the DSP board, a 1024 point FFT on complex data can be done in 3 msec using a C callable assembly language routine. To do the same algorithm on the 486 CPU requires much more time.

The DSP board also operates independently of the CPU and has memory that can be accessed by the CPU without stopping the board processing. This allows the CPU to control the A/D cards, transfer data and do some processing while the DSP board does the signal processing. The processing time is less than half of what it would be if everything was done by the CPU.

The azimuth compression is not done in real-time because it requires collecting 512 range-compressed pulses before azimuth compression can be done. For azimuth compression, data is stored in a stream of range compressed segments, but must be accessed by the range location. To avoid slow hard drive accesses, the data is stored on a RAM drive. Azimuth compression is also slowed by the requirement to generate a new azimuth chirp and transform it to the frequency domain for each range bin. This adds another FFT to each processing step as well as the calculations required to create the chirp.

The range and azimuth compressed data is converted to a magnitude format for display. A magnitude format is achieved by taking the square root of the sum of the squared real and imaginary parts for each data point. Using magnitude instead of power data samples shows more detail because the ratio of a small to big number decreases by squaring them. The data are scaled into the range 0 to

63 (only a 6 bit display is available), converted to byte format and displayed on the PC monitor as a grey scale image. It is also stored as a byte file with data values ranging from 0 to 255 for printing.

4.5 System Cost, Performance and Physical Description

In this section, the BYU SAR is compared to other current SAR systems. Most U.S. SAR systems are owned by the military, but since specifications are classified and use denied, BYU SAR is compared to privately owned U.S. SAR systems. In the U.S. the two operating private SAR systems are the Jet Propulsion Laboratory (JPL) SAR and the Environmental Research Institute of Michigan (ERIM) SAR. The system parameters for BYU, JPL and ERIM SARs are compared in Table 1.

It can be seen that the goal was reached in obtaining a low cost SAR system with good performance. The cost of the lower power BYU SAR is about a hundredth of the JPL and ERIM SARs' costs. The chief issue in performace is resolution which is quoted in Table 1, but demonstrated in Chapter 5. The ERIM SAR has better range resolution, but they use a rectangular window on their transmitted chirp so their sidelobes are only 13 dB down from the main lobe and use 6 bit quantization.

Parameter	BYU SAR	ERIM	JPL
Cost $	$ 70 K	in the $ 10 M range	in the $ 10 M range
Maximum Power	1 watt		1 KW, 6 KW, 1 KW
Processing	near real-time	near real-time	real-time
Range resolution (δ_r)	2.5 m	1.5 m	7.5 m, 3.75m
Azimuth resolution (δ_{az})	0.45 m	1.44 m	8 m (4 looks), 2 m
Carrier frequency (f)	10 GHz	9.35 / 1.185 GHz	.44, 1.25, 5.3 GHz
Chirp bandwidth (B)	90 MHz	90, 60 MHz	20, 40 MHz
Chirp length (τ)	1.6 μsec	2.7 / 1.8 μsec	11.25 μsec
Pulse repetition frequency (PRF)	239 hz		250 - 750 hz
antenna polarization	vv	quad	quad
Antenna beamwidth	3°		19°, 8°, 2.5°
quantization	8 bits/sample	6 bits/sample	8 bits/sample

Table 1: System parameters for the BYU, ERIM and JPL SARs.

The BYU SAR's resolution in range and azimuth is better than the JPL SAR's with the same quantization so the BYU SAR is at least as good as the other two systems in resolution and quantization.

A characteristic of the SAR systems that is not listed in Table 1 is the size of the systems. The BYU SAR is compact and light enough to be deployed from a variety of platforms such as a Cessna size airplane, but for testing it was deployed from the back of a one-ton truck. The other systems are deployed from a D-8 size airplane. For vehicle based operation of BYU SAR, three pieces are secured for travel as shown in Figure 13: a power generator, the antennas and an equipment rack. The truck and antenna setup is shown in Figures 14 and a diagram of the rack is shown in Figure 15. This is a prototype setup and no attempt was made to minimize the size; it could be made much smaller.

Figure 13: Photograph of deployment setup showing 3 deployment pieces.

The BYU SAR does not look as versatile as the JPL and ERIM SARs in that it only has one carrier frequency, but this is only a limitation of the antenna. With different antennas, a different carrier frequency could be used. A different polarization could be chosen for any given deployment as well, but not simultaneously.

The platform used for the antennas is subject to vibration and swaying. This changes the return data because a change in the antenna direction corresponds to a change in the azimuth waveform. This causes a problem if the azimuth chirp is changing rapidly over the processing length of the data. The

Figure 14: Photograph of deployment setup showing truck and antennas.

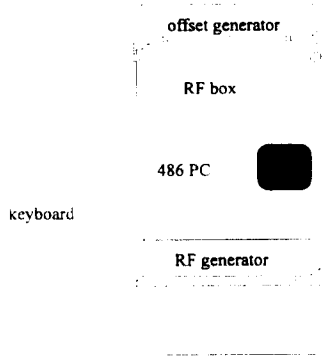

Figure 15: A diagram of the rack.

antennas that were used had a high gain and a narrow beamwidth which corresponds to a narrow bandwidth azimuth chirp. This makes the system very susceptible to the errors caused by vibration and swaying. The antenna beam can easily move far enough to shift the azimuth chirp completely out of the frequency pass band of the correlating chirp. To help with this, the person riding in the back of the truck not only runs the software, but also manually stabilizes the antennas. In spite of the difficulties, we were able to successfully image a site, the results of which are detailed in the next chapter.

Chapter 5

SAR Performance and Test Image

The BYU SAR, while only a prototype, attained high performance. The design goals included achieving 1.5 meter resolution in the azimuth dimension and 2.5 meter resolution in the range dimensions, each with reasonably low sidelobe interference. The resolution was demonstrated by imaging a site on the west side of University Avenue near BYU in northern Provo Utah. The deployment specifics are given in section 5.1. Section 5.2 compares the system performance with the design criteria by examining strips from the image and looking at the range and azimuth resolution. It is shown that the design criteria is met in the range dimension and exceeded in the azimuth dimension. Section 5.3 discusses the University Avenue image by features and the ability of the BYU SAR to map a site is examined. It is seen that many man-made as well as natural objects in the field are identifiable in the image, demonstrating the usefulness of the BYU SAR as a remote sensing tool with applications in archaeology.

5.1 Deployment Overview

The BYU SAR was deployed from the back of a one-ton flatbed truck (see Figure 13) at a speed of 35 mph on University Avenue. The area imaged is just north of the BYU livestock area located on 2230 North in Provo. The site was chosen because University avenue is straight and smooth and the site has a wide range of useful targets. The site has point targets that are good for determining resolution and distributed targets that help demonstrate the SAR's imaging ability. Targets vary in range from two to 150 meters so the SAR's success is demonstrated at different ranges. There is also a combination of man-made and natural targets that help demonstrate BYU SAR's usefulness in archaeology.

5.2 Image Quality

The two criterion that will be used to judge the quality of the image produced by the BYU SAR are resolution and sidelobe levels. The resolution is a measure of how close two similar objects can be and still appear as two targets. This can be measured by the size of a point target's response in an image. The resolution was about 3 meters in the range direction and 0.45 meters in the azimuth direction. The sidelobe level is a measure of the height of the sidelobes of the compression function. High sidelobes can cause artifacts and make the image appear more noisy. The resolution and sidelobe levels were measured by extracting strips with bright point like targets from the University avenue image. The range resolution was measured at one range distance and the azimuth resolution was measured at two different ranges.

5.2.1 Image Quality in the Range Dimension

The achieved range resolution was equal to the ideal resolution defined by the transmitted chirp. In Chapter 3 we showed the time resolution after pulse compression is equal to the inverse of the chirp's bandwidth for a chirp with a high time bandwidth product. The chirp bandwidth of 90 MHz corresponds to a time resolution of 22.2 nsec. The radar wave travels at approximately 3×10^8 m/s which gives a distance of 6.7 meters which corresponds to a null to null range resolution of 3.3 meters and a 3 dB resolution width of 1.5 meters. A raised cosine window was used on the chirp to reduce sidelobes in the frequency domain after correlation. This causes the resolution cell size to be increased by a factor of 1.5. Thus, the anticipated resolution is 5 meters at the nulls and 2.2 for the 3 dB points.

Figure 16 is a strip in the range dimension extracted from Figure 19 showing that the achieved range resolution is equal to the ideal resolution for the transmitted chirp. The strip contains a fencepost in the range dimension taken from the cross-wire fence in the middle of the image. The dashed line represents the ideal resolution of a point target dictated by the transmitted chirp. The range resolution in the image is just a quantized version of the ideal resolution. The waveform has a maximum value of 145 so 3 dB down is 100. The resolution shown is about 2.5 meters in the range direction, but is limited by the 1.5 meter pixel size. The first sidelobes are down below the qantization level in the theoretical plot and appear to be the same in the strip showing the actual data.

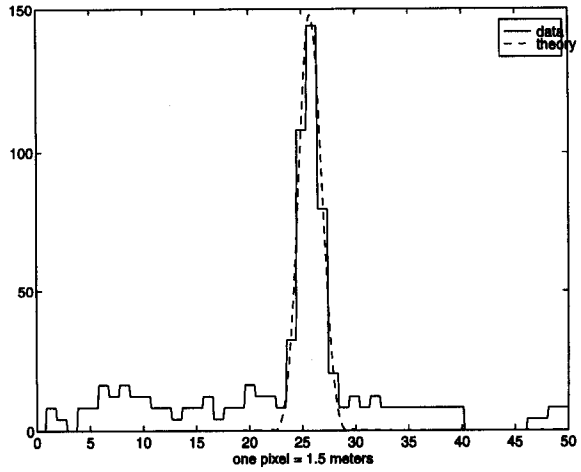

Figure 16: A strip showing the range resolution.

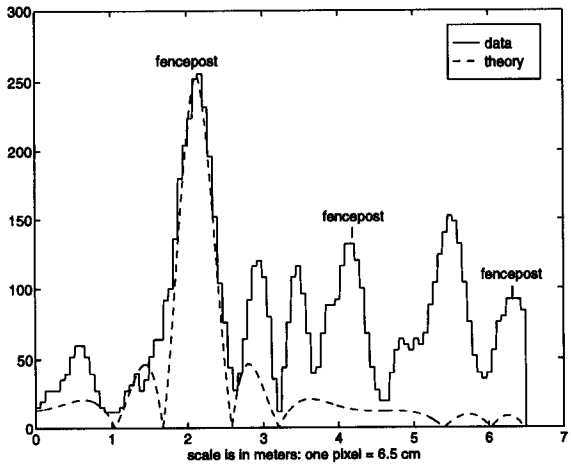

Figure 17: A strip from cross-wire fence showing azimuth resolution. The dashed waveform is the theoretical resolution.

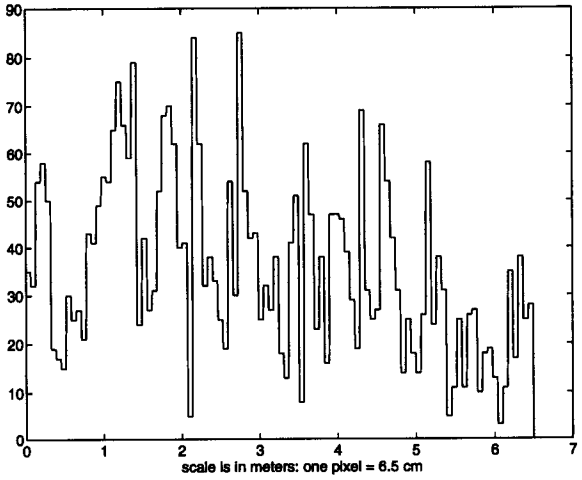

Figure 18: The uncompressed strip used in Figure 17.

5.2.2 Image Quality in the Azimuth Dimension

The SAR achieved azimuth resolution that is better than the design criteria. The design criteria was 1.5 meters, but targets at 90 meters and 120 meters generated images of 0.45 meter resolution in

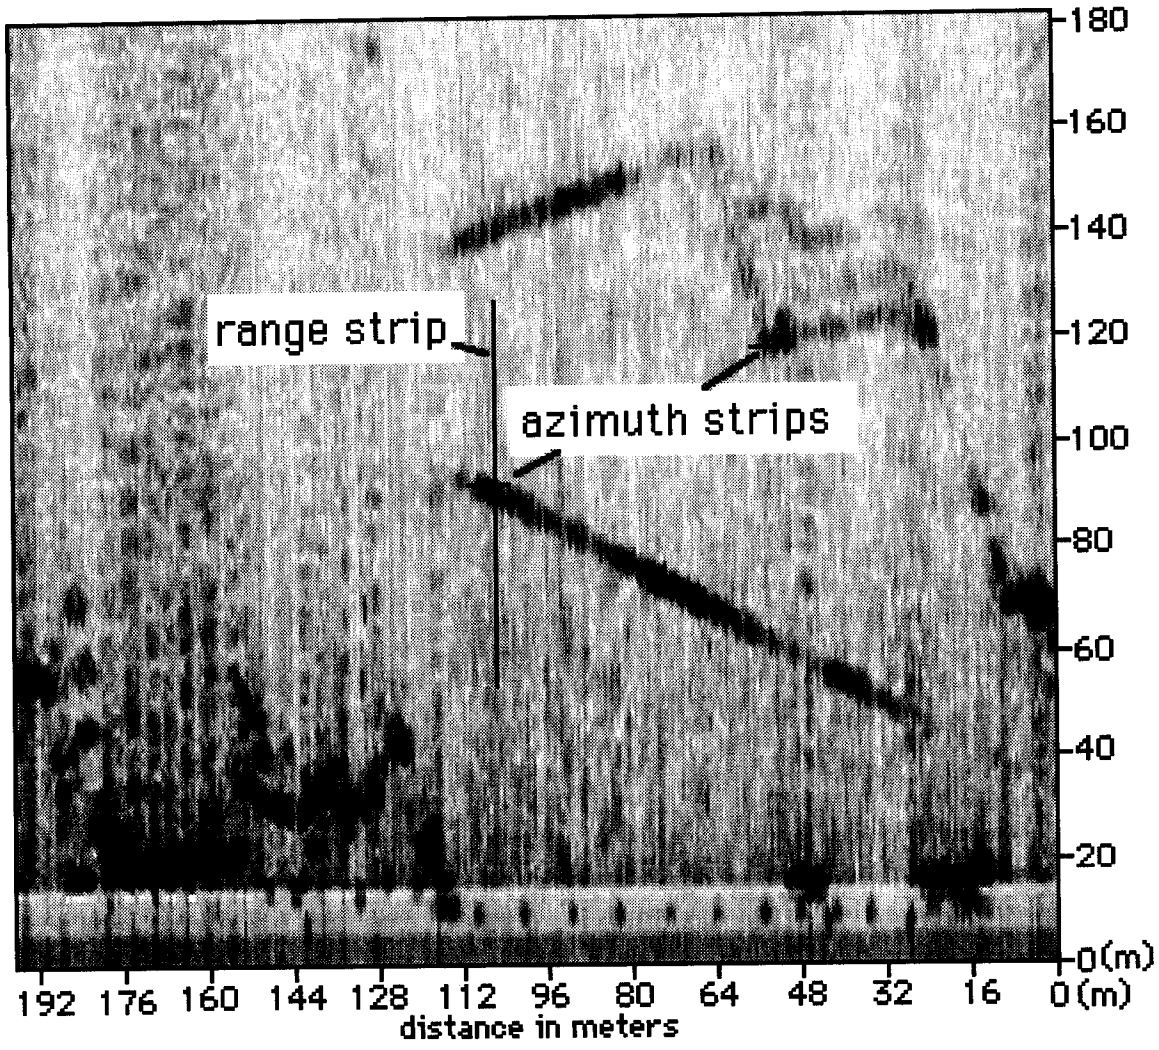

Figure 19: Image showing where the range and azimuth strips were taken from.

the azimuth direction. Descriptions of the appropriate parts of the Universtiy Avenue image are presented to illustrate the high resolution.

Figure 17 shows the quality in the azimuth dimension of the image at 90 meters. The expected response for a point target is a sinc function with the nulls of the main lobe about one meter apart. Figure 17 is a strip showing fence posts that act like point targets on a cross-wire fence which is a distributed background. The resolution is shown by the response from the first fencepost since it is bright enough to stand out from the wire mesh response. It is about seven pixels or .45 meters wide at the 3 dB point (170). This can be compared to Figure 18 which is the same strip compressed in range, but not in azimuth. The two figures show that the azimuth compression is working and the azimuth resolution is almost exactly equal to theory.

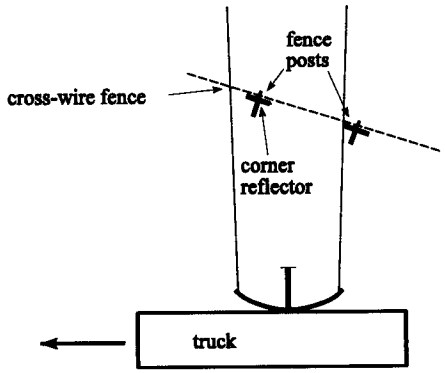

Figure 20: Fence post orientation.

The fence posts were separated by two meters in the azimuth dimension so the next fenceposts are labeled accordingly. The main lobe for this fencepost is less bright. This could be caused by speckle noise or a different post orientation. The fenceposts had a vertical cross section that looks like a letter "T" with the bottom of the letter pointing at the road as shown in Figure 20. This makes two corner reflectors for the vertically polarized radar wave. Its brightness would be affected by how much and in what direction it is leaning as well as the direction the letter "T" is pointing. It is about 60 cm closer in range than the main post, hence, the strip cuts through the side of the second post's response.

The sidelobes are a little higher than expected, but this could be caused by a slight inaccuracy in the synthesized chirp. The azimuth correlation chirp is windowed with a rect function, but the data is windowed with the two way antenna beam pattern which isn't known exactly. The azimuth windowing can be approximated by a rect function since only the data between the 3 dB points is used. The sidelobe amplitude should be only slightly less than that of a rect window (13 dB down from the mainlobe) because of this. The peak on the right has a magnitude of about 13 dB so the windowing by the antenna

beam was offset by an inaccuracy in the chirp rate.

The peaks on the left are higher which means that more of the vertical strips in the fence are showing up in addition to the fenceposts. These peaks are also being combined with a sidelobe from the post at 2.4 m, but they are much narrower than the sidelobe on the right. This could be because the fence is more distributed making the response less constant due to interference and speckle noise. These peaks also have the same width as the top of the mainlobe for the fence post at 4.5 m. This suggests that they could be ghost images. The actual image supports this since the peak of the ghost image is at the same range bin as the peak to the right.

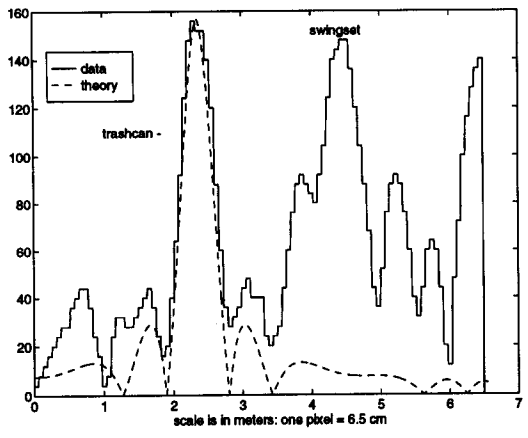

Figure 21: A strip showing azimuth compression @ 120m the dashed line shows the theoretical resolution.

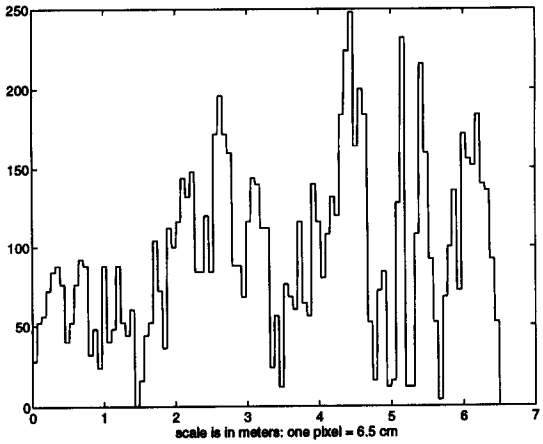

Figure 22: The same strip as shown in Figure 21 without compression.

Figure 21 is a compressed azimuth strip from 120 meters that shows some of the same characteristics as Figure 17. The expected resolution, for travel at 35 mph and a 3° beamwidth, is shown with a dashed line. The BYU SAR shows the optimal azimuth resolution for the system parameters (speed, beamwidth and PRF). Figure 22 is the same strip without compression. The peak labeled "trashcan" is the return from a 30 cm wide metal trash can. It was a little flattened so it has a brighter response than a round can. The beam resolution for this range appears to be about the same as at 90 meters, 7 pixels or .45 meters. This means that the time-bandwidth product at both ranges is high enough for the algorithm to be successful at both ranges. The two peaks to the left and the peak just to the right are 13 dB down from the peak amplitude and have the right width to be sidelobes. The swing set that is labeled in the strip is 1.6 meters wide at the bottom. It appears to have ghost-like images to its right. These are probably from the fence directly in front of the swing set which was a thick metal

grid with six inch spaces between the vertical strips.

An explanation for the extremely peaked appearance of the fence is sidelobe cancellation. This is a result of the negative sidelobes of the correlation sinc functions cancelling a portion of the main lobe of a nearby target. A simulation is done on a combination of three targets that each have some azimuth spread. Figure 23 shows the slightly distributed targets while Figure 24 is the correlation of the synthetic targets with the autocorrelation of a chirp. The azimuth spread causes the splitting of the main lobe as is seen in the splitting in the main response at 4 m of Figure 21. This indicates that some of the small peaks in the split main lobe that look like artifacts may be caused by sidelobe cancellation from other targets or the distributive nature of a target.

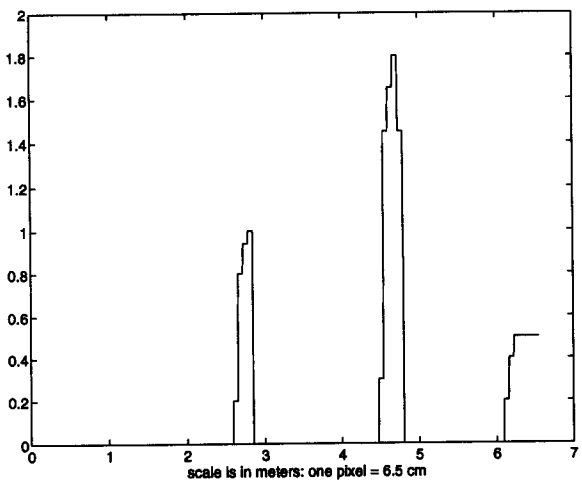

Figure 23: Synthetic targets for demonstration of sidelobe cancellation.

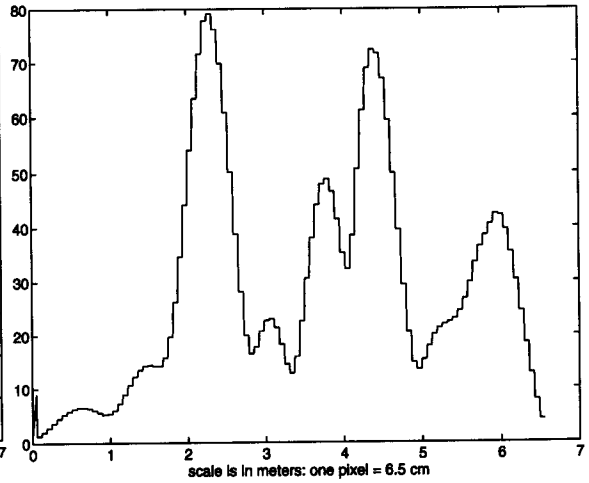

Figure 24: Correlation of targets demonstrating sidelobe interference.

5.3 University Avenue Test Image

In this section the University Avenue image is examined one feature at a time. Figure 25 is the image of the site at a speed of 35 mph with no compression. Its range axis is longer to show all of the uncompressed waveform. Figure 26 only has range compression and shows that the success of the range compression is consistent throughout the entire image. A lot of detail can be picked out even in the foreground of the image. The range scale is scaled so that the image is proportional with 160 meters in range and 150 meters in azimuth. Figure 27 shows that both range and azimuth compression work and targets are much sharper in appearance after azimuth compression. Figure 28 is a diagram of the site east of University avenue.

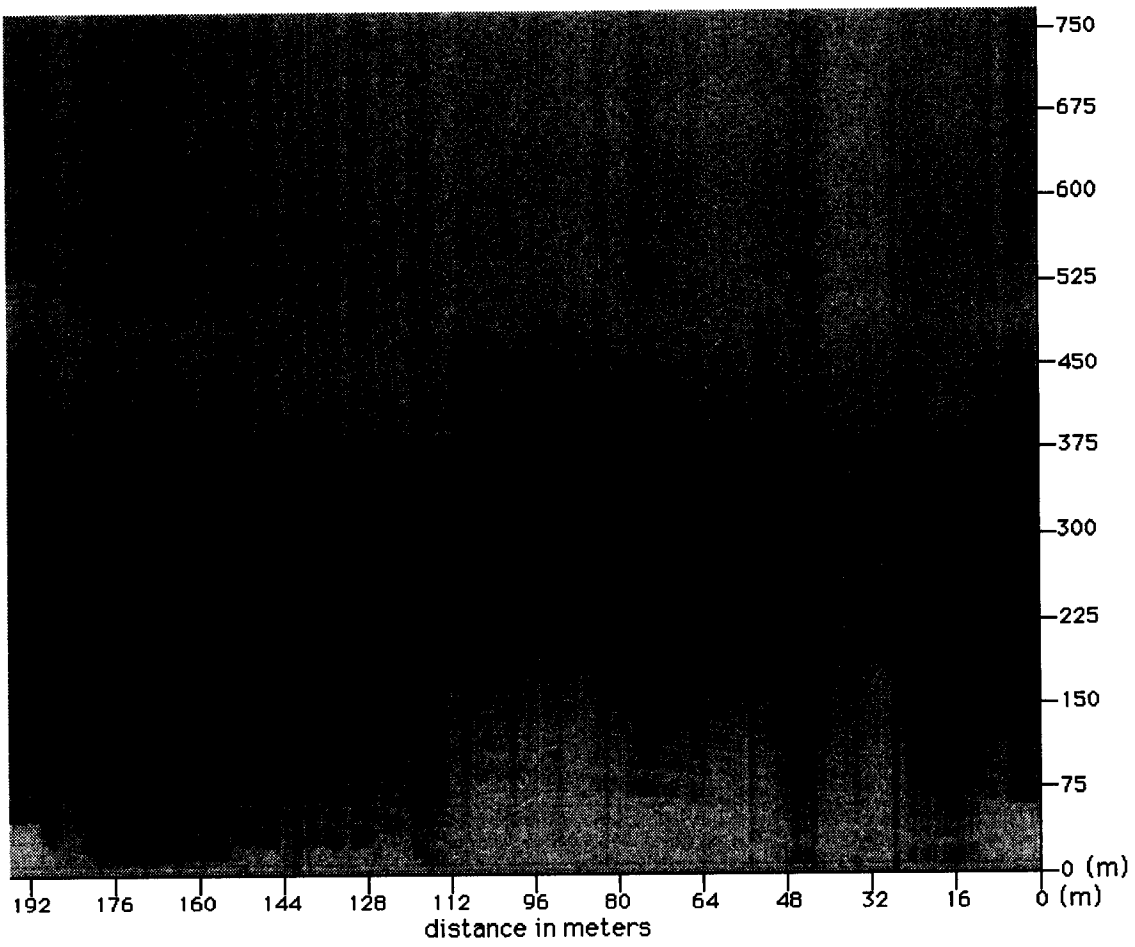

Figure 25: Uncompressed image.

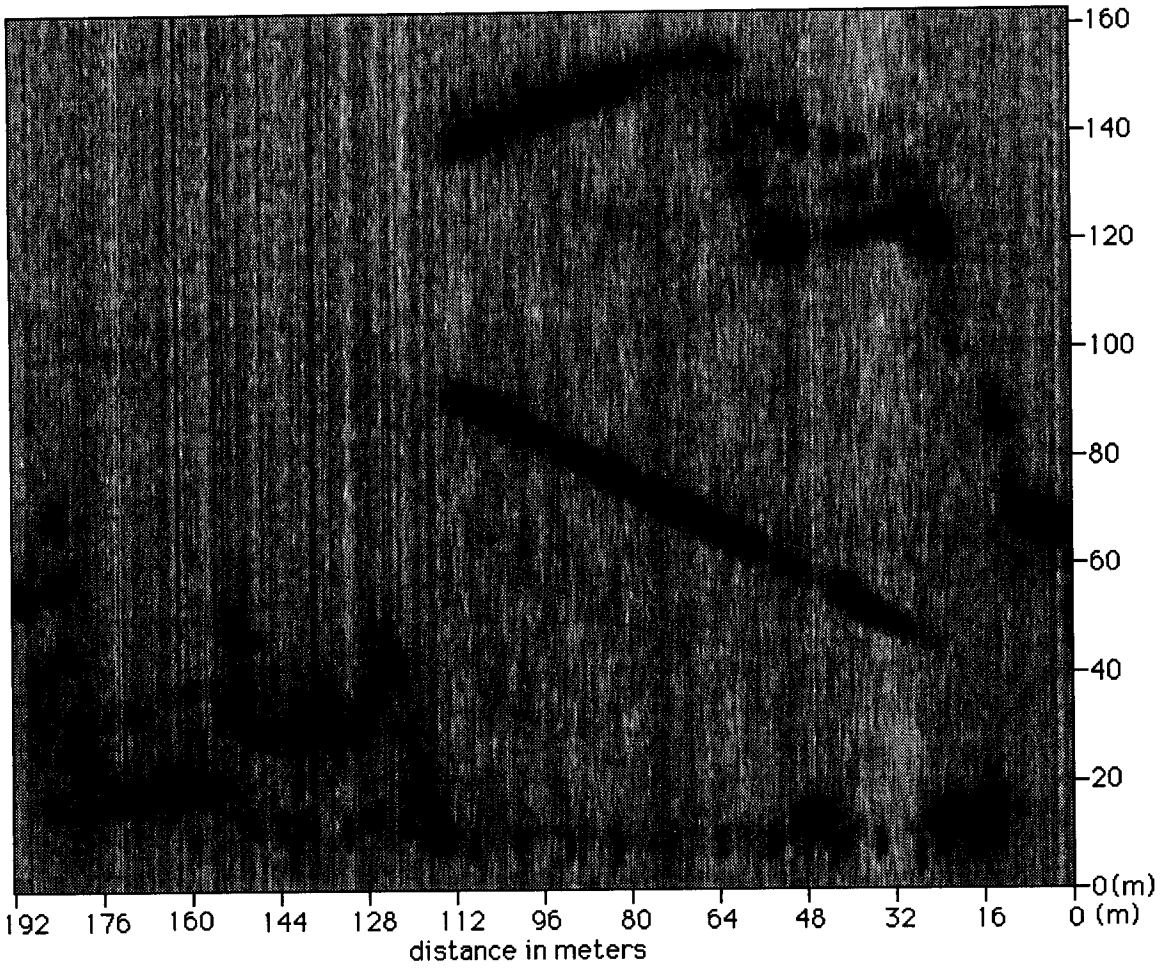

Figure 26: Range compressed image.

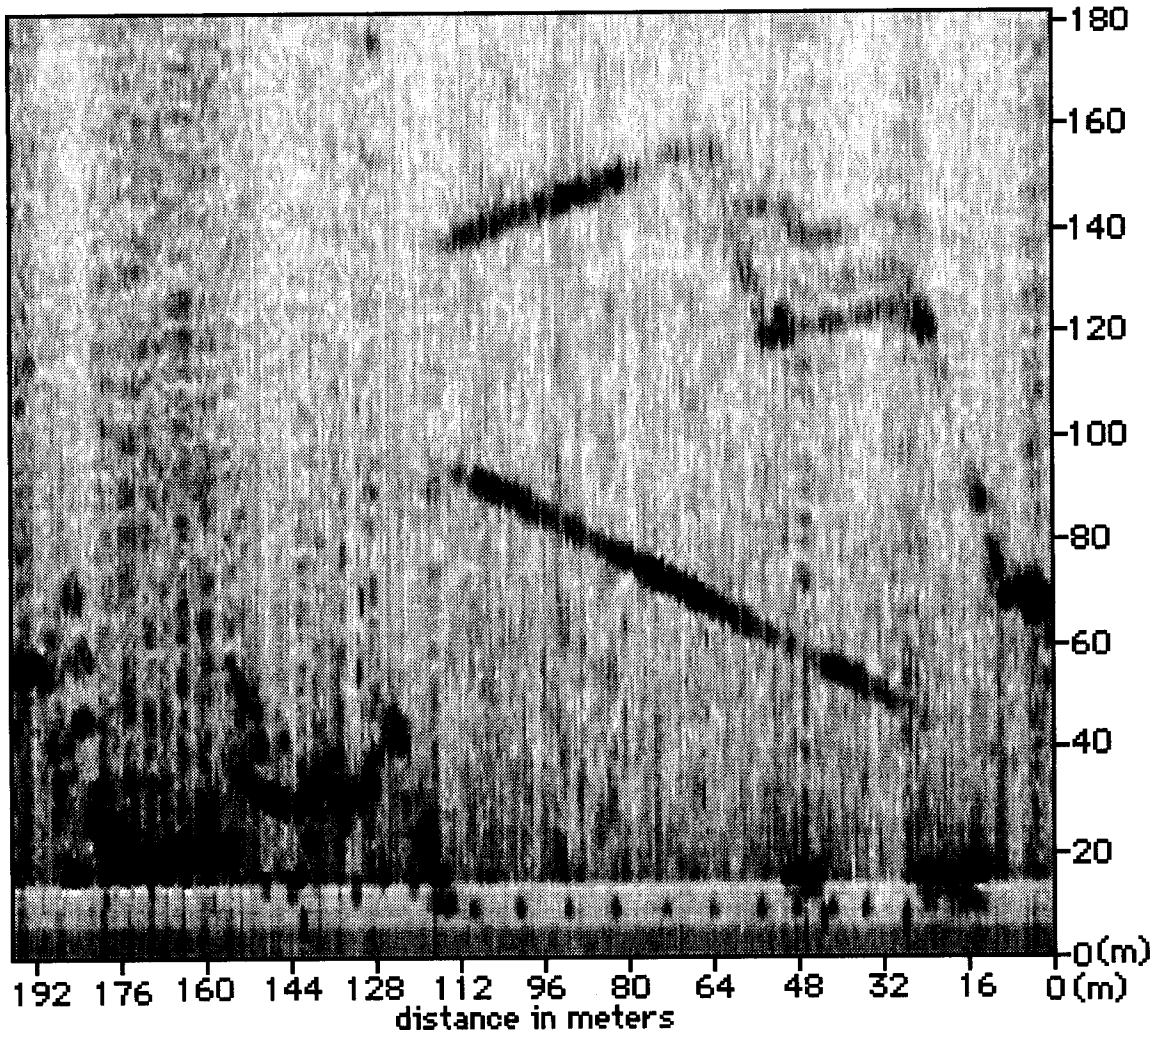

Figure 27: Fully compressed image.

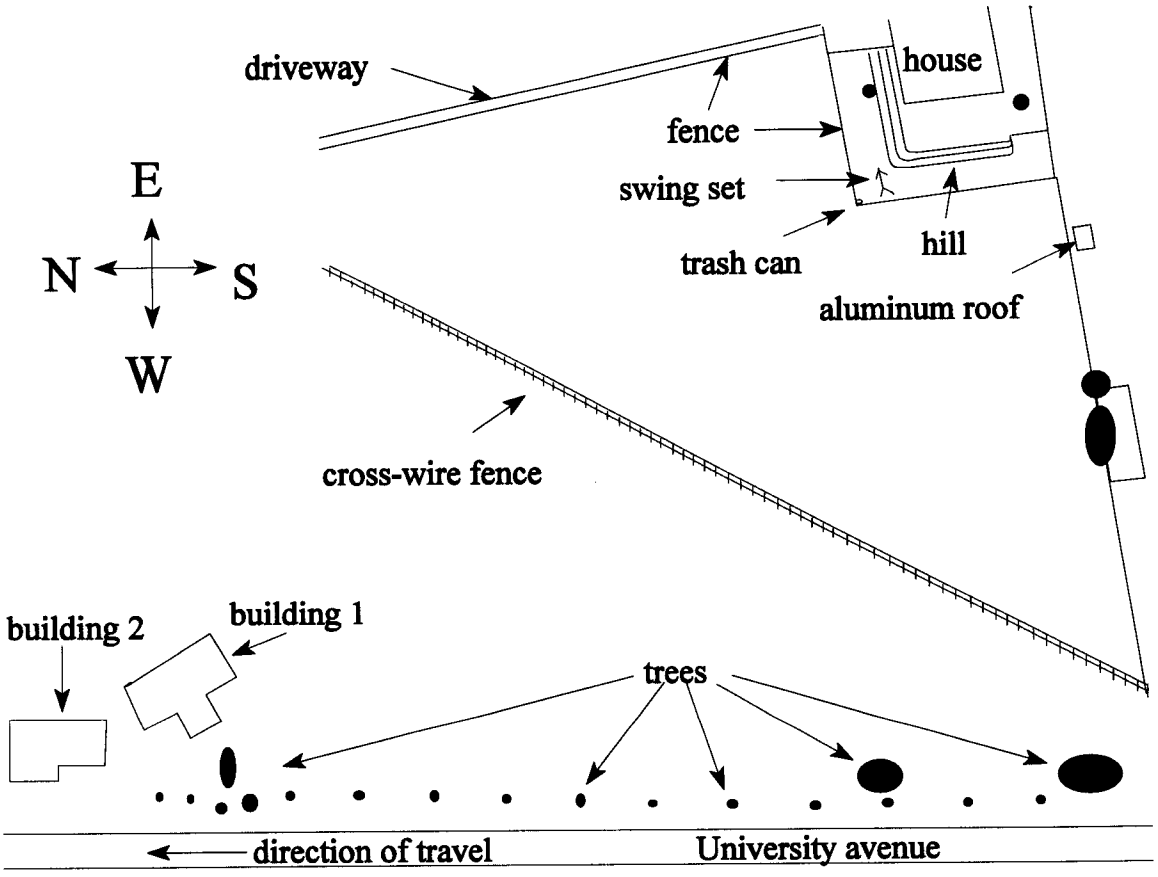

Figure 28: A diagram of the imaged site.

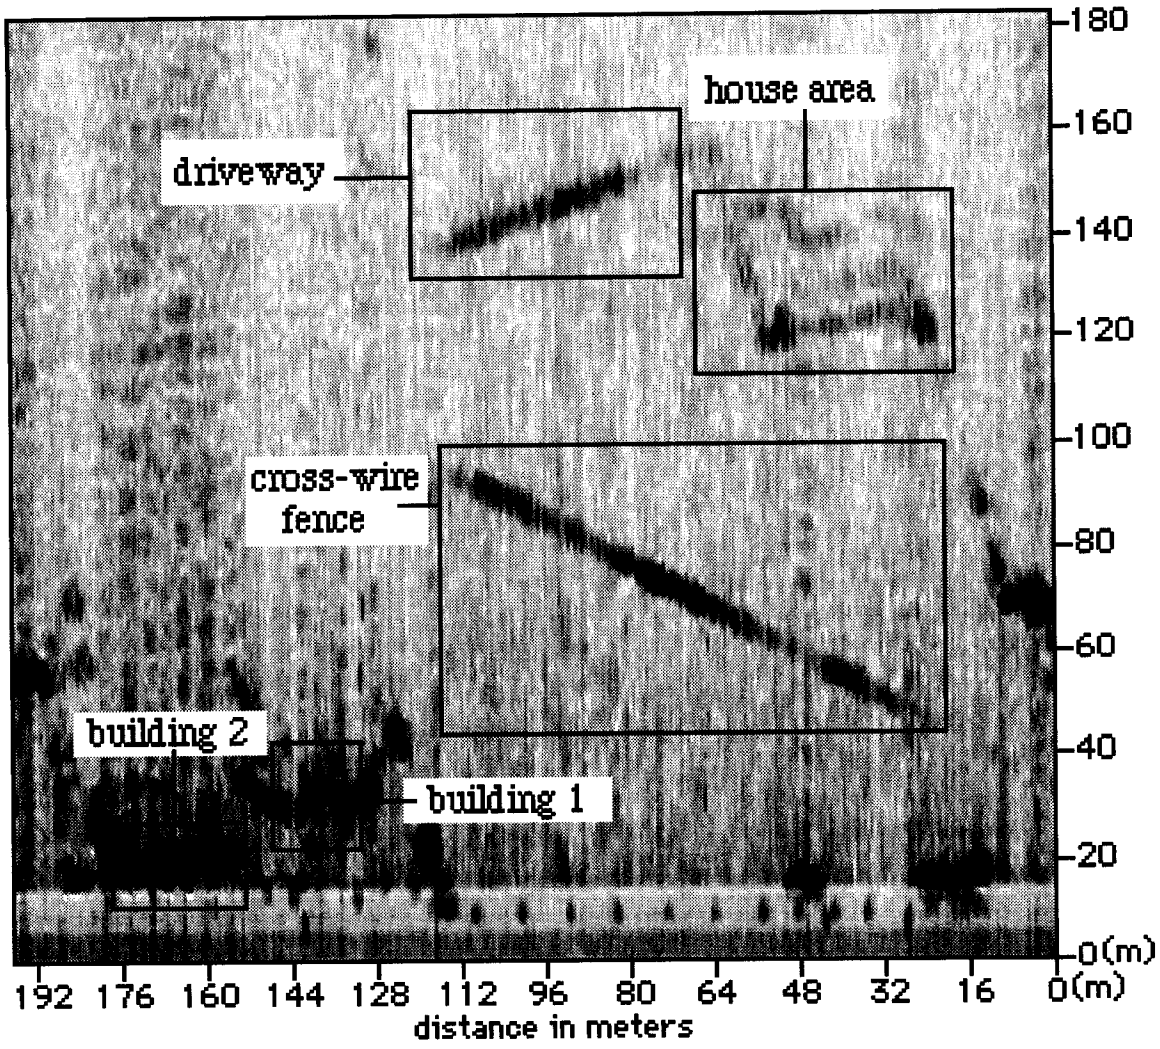

Figure 29: Significant features.

The University Avenue image is discussed according to the features highlighted in Figure 29. Section 5.3.1 references the features that can be seen in the foreground of the image with good range resolution and a narrow antenna footprint. Section 5.3.2 discusses the detail in the cross-wire fence and what insight it gives as to what is happening in acquiring the data. Section 5.3.3 describes the objects that can be identified in the image that are located around the house. It also makes the point that without the house, the living area would still be identifiable for archaeologists in the distant future. The last section makes a similar point about the driveway at the end of the field.

5.3.1 Foreground

The foreground of the image or the portion of the image that is close in range has good image quality. In this section the SAR image is similar to what would be received from a Side Looking Aperture Radar or SLAR. A SLAR does range compression and relies on the antenna beamwidth in the azimuth dimension to determine the azimuth resolution. The antenna beamwidth on this deployment was narrower than the range resolution for the first 17 meters of range so even though the time bandwidth product is very small, the azimuth resolution is still very good. The range compression works very well for the entire image so the image quality in both the foreground and the background is good.

A lot of detail can be seen in the foreground of the image. Figure 30 is a photograph of building 1 shown in Figures 28 and 29 that appears in the foreground of the image. The outline of the front of the building is a very bright target in the image. Building 2 in Figures 28 and 29 is also outlined clearly in the image. The small evenly spaced dark spots on the image foreground are trees similar to the small trees in Figure 30. The resolution is demonstrated by resolving the road sign and a small tree in the south side of the foreground. The large bright targets in the evenly spaced trees are just bigger trees. Even though the SAR algorithm is not as effective in the foreground of the image, with good range resolution and a narrow antenna footprint, objects in the foreground of the site are identifiable.

5.3.2 The Diagonal Fence

The diagonal fence illustrates how well the SAR is working. The fence is the cross-wire fence labeled in Figures 28 and 29. It has a 15 cm separation in the wires of the grid. It also appears in the image in Figure 27 and in the photograph in Figure 31. The vertical strips in the mesh added to the strong response from the fenceposts. The azimuth resolution helped to make the fence a narrower structure in the image. It also brought out the fenceposts individually. The separation of the fenceposts appears nonuniform along some portions of the fence. This is caused by the horizontal swinging

oscillation of the antenna. The Doppler frequency is mapped to position in the azimuth dimension and so when the Doppler centroid changes without using a different compression chirp, a spatial bias occurs.

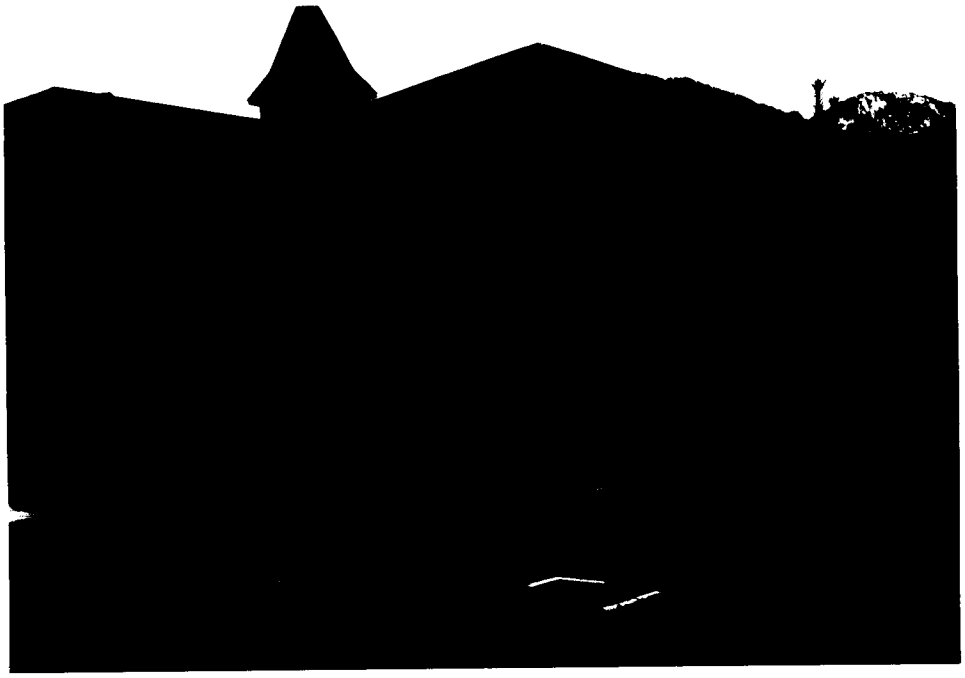

Figure 30: Photograph of building 1.

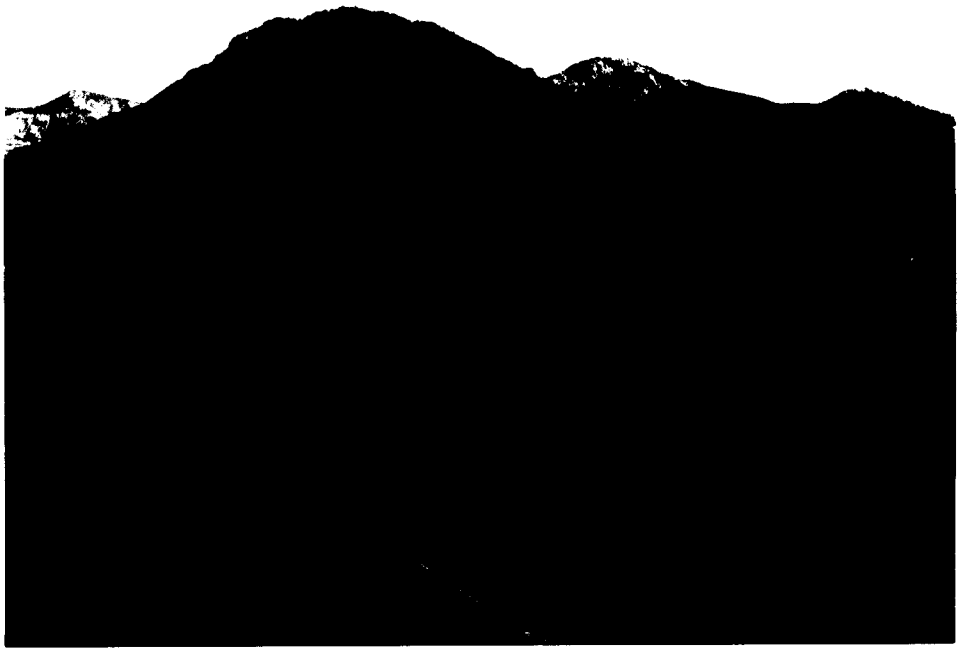

Figure 31: Photograph of site showing diagonal fence and house area.

5.3.3 The House Area

The house area is rich in both man made and natural targets that are typical for a living area. The fence in front of the house was a bright target with two sections that had thicker wire and therefore, a brighter response than the rest of the fence. In addition to the fence, a small metal trash can was in the corner of the yard west of the house. On the side of the fence closest to the edge of the image is a shed with an aluminum roof that was very bright. The house is an example of a man made structure showing up well in a SAR image. The house is elevated and the antenna pointed slightly down. Therefore it was in an attenuated portion of the antenna beam.

The hill the house is built upon is also easy to see in the image. This demonstrates the SAR's ability to resolve ground relief. It is interesting to note that if the house was gone, the SAR image would still show the foundation as evidence that the house was previously there. This would be very useful in archeology when the structures of a society are no longer standing. The other bright targets in front of the hill are a swing set, a trash can and a portion of the fence with thicker wire. These demonstrated the good azimuth resolution of the SAR at greater range distances.

5.3.4 The Driveway and the South Fence Line

The driveway behind the house is another very bright target. It has a fence in front of it and a hill with several trees and bushes behind it. It is very bright, but there are many targets that contribute to its brightness which also makes it very susceptible to speckle noise. It is also a permanent indication of a home and could be identified without the presence of the buildings. The only other bright target that appeared in the image is the tree and building that are on the south side of the image.

5.4 Summary

We have seen that the BYU SAR demonstrates good performance. The resolution was 2.5 meters in the range direction and 0.45 meters in the azimuth direction which meets the design criteria. The BYU SAR also demonstrated the ability to map a site by imaging a field in northern Provo with many targets of interest in a well scaled image. In addition to the point like targets that were used to demonstrate resolution, ground relief was identified and many man-made targets are easily seen in the image. This demonstrates a usefulness in archaeology.

Chapter 6

Conclusions

6.1 Summary of Results

In this thesis, current technology and an innovative design were used to produce a low cost SAR with competitive performance. The cost of the system is about one hundredth the cost of current systems making SAR accessible to new fields of study. The main criteria of performace, the resolution of the SAR, is competitive with the privately owned SAR systems in operation today. The system demonstrates effective site imaging and a usefulness in archaeology.

The BYU SAR makes a SAR system accessible for more applications by providing a low cost SAR. The need for an IF frequency is eliminated by an on-board processor. The use of improved technology in arbitrary waveform generation and sampling have helped to reduce the cost by performing all of the baseband processing digitally. An innovative design was employed to allow the transmission of a real waveform to save additional cost. As a result, the cost of BYU SAR is around $70,000.

The resolution of the SAR system in this thesis is competitive with current private systems. The range dimension resolution is 2.5 meters with very good sidelobe suppression. The azimuth resolution is 0.45 meters, but with higher sidelobes. This is competitive with the highest resolution of unclassified SAR systems.

The BYU SAR is effective in site imaging. Many features were identifiable in the image of a field in northern Provo. The outline of a house is identifiable, in addition to the house's foundation. The foundation is evidence of the house even when the house is gone and is therefore an example of the usefulness of SAR in archaeology. With its resolution and cost, BYU SAR is an answer to Irwin Schollar's criticism of the use of SAR in archaeology, saying that the cost was too high and the resolution was insufficient.

6.2 Improvements on the Current System

The current system meets the design criteria, but with some modifications of the system, BYU SAR can have higher performance and be useful in a broader range of applications. The system size could easily be reduced to allow deployment from a small vehicle or airplane. The A/D cards were another shortcoming of the system; data reliability could be improved with a different A/D system. Speeding up the code and using a more sturdy antenna mount are two other ways to improve the system.

This prototype deployment setup served its purpose in demonstrating a working SAR, but all three deployment items (the rack, power source and antenna mount) could be made much smaller. Much of the equipment used in testing the SAR system was just material that was on hand. Batteries could supply the system power instead of the gasoline power generator and alternative function generators could be used that are built specifically for generating signals at the desired frequencies.

The mount for the antennas was salvaged from another project and a new mount could improve the system. The antenna mount allowed steering of the antennas but the antennas could be secured to a much smaller structure. The antenna platform vibration was a big source of data corruption and could be eliminated. A platform that is sturdy and doesn't allow large amplitude physical oscillations of the antennas in the horizontal direction from wind and acceleration would improve the quality of data. If the oscillations in this direction were eliminated, the azimuth compression algorithm would be much more successful. A well structured wood stand would be effective and low cost.

The A/D cards corrupted data by the uncertainty in the trigger address. The trigger address uncertainty ruined the I/Q structure of much of the data and added an uncertainty in the range direction. The cards that were used were the best advertised at the time, but since then, cards with higher performance have become available. Different A/D cards could also make it possible to achieve a higher PRF.

Increasing the speed of the computer code would allow a higher PRF for real-time operation. Assembly language code could be used in multiplting the Fourier transform of the chip with the transform of the data. In addition to this, it might be faster to transfer the data directly from the A/D cards to the DSP board.

6.3 An L-Band System

The SAR system as it is now exists is a prototype for one that can be used in archaeology. The carrier frequency was chosen to be 10 GHz so that available antennas could be used for testing, but as

was stated in Chapter 2, a lower carrier frequency is desirable for archaeological applications. This would allow penetration of vegetation canopy and dry sandy soils. The RF circuit board can be used at frequencies as low as 2 GHz, and so at that frequency only new antennas would be required. The same design could be used in manufacturing a new circuit board with RF components that function at the lower frequencies for a less expensive board if the system design was for lower than 2 GHz. A higher power amplifier could be used with an L-band system that would give an output power of 10 watts for up to 900 MHz at half the cost of the 1 watt amplifier in the current X-band system.

The antennas need to be designed with the application in mind. The maximum Doppler shift from a target determines the PRF which may be a limiting factor. The Doppler shift is proportional to the wavelength so using a lower carrier frequency will create a lower Doppler shift. This allows for more freedom in the required PRF and antenna beamwidth. For a real-time implementation, the maximum PRF is 31.8 Hz. Sampling in azimuth at 31.8 Hz for a 0.9 GHz system and allowing an operating speed of 50 mph, limits the maximum beamwidth in the azimuth dimension to 27°. This means that the antenna must be at least .7 m or a little more than two feet in the azimuth dimension. An antenna this size would be relatively easy to deploy, but its gain would be much lower than the antennas presently being used. This might not be a problem since a 10 watt power amplifier would be available. In addition, the RF board could give an additional 10 dB of gain on the receive line by eliminating a 10 dB attenuator that is used to avoid waveform clipping.

Appendix

A.1 The Matched Filter Proof

This proof of the matched filter is adapted from Fitch (1988). The matched filter is the filter that gives the optimum SNR for a deterministic signal with additive noise. The signal to noise ratio is given by

$$SNR(t)=\frac{S_0^2(t)}{E\{N_0^2(t)\}} \tag{13}$$

where $E\{\cdot\}$ is the expected value. The signal to noise ratio is a maximum at t =T. $S_0(t)$ is the filtered signal s(t) and N_0 is the filtered noise $\eta(t)$ given by

$$S_0(t)=\int_{-\infty}^{T} s(\tau)h(T-\tau)d\tau \tag{14}$$

and

$$N_0(t)=\int_{-\infty}^{T}\eta(T-\tau)d\tau. \tag{15}$$

Then the signal to noise ratio can be expressed as

$$SNR(T)=\frac{\left|\int_{-\infty}^{T} s(\tau)h(T-\tau)d\tau\right|^2}{\int_{-\infty}^{T}\int_{-\infty}^{T} R_{\eta\eta}(T,\sigma)h(T-\tau)h(T-\tau)d\tau\,d\sigma} \tag{16}$$

where $R_{\eta\eta}(\tau,\sigma)$ is the autocorrelation of the noise $\eta(t)$ given by $E\{\eta(\tau)\eta(\sigma)\}$. Next, the maximum signal to noise ratio is defined as $1/\alpha$ so that we have

$$E\{N_0^2(t)\}-\alpha S_0^2(t)\geq 0 \tag{17}$$

or

$$\int_{-\infty}^{T}\int_{-\infty}^{T} R_{\eta\eta}(\tau,\sigma)h(T-\tau)h(T-\tau)d\tau\,d\sigma - \alpha\left|\int_{-\infty}^{T} s(\tau)h(T-\tau)d\tau\right|^2 \geq 0. \tag{18}$$

The maximum is given when the equality holds. To solve this, calculus of variations is used.

$$\int_{-\infty}^{T}\int_{-\infty}^{T} R_{\eta\eta}(\tau,\sigma)[h(T-\tau)+\epsilon\delta h(T-\tau)][h(T-\tau)+\epsilon\delta h(T-\sigma)]d\tau\,d\sigma - \alpha\left|\int_{-\infty}^{T} s(\tau)[h(T-\tau)+\epsilon\delta h(T-\tau)]d\tau\right|^2 = 0 \tag{19}$$

This is a polynomial in epsilon of the form $P(\epsilon) = P_0 + 2\epsilon P_1 + \epsilon^2 P_2$ where P_1 is given by

$$P_1 = \frac{1}{2}\int_{-\infty}^{T}\delta h(T-\tau)\int_{-\infty}^{T} R_{\eta\eta}(\tau,\sigma)h(T-\sigma)d\sigma\,d\tau + \frac{1}{2}\int_{-\infty}^{T}\delta h(T-\sigma)\int_{-\infty}^{T} R_{\eta\eta}(tua,\sigma)h(T-\tau)d\tau\,d\sigma - \frac{\alpha}{2}\int_{-\infty}^{T} s(\sigma)h(T-\tau)d\sigma\int_{-\infty}^{T} s(\tau)\delta h(T-\tau)d\tau - \frac{\alpha}{2}\int_{-\infty}^{T} s(\tau)h(T-\tau)d\tau\int_{-\infty}^{T} s(\sigma)\delta h(T-\sigma)d\sigma.$$

Extrema of the polynomial $P(\epsilon)$ can be found by noticing that $\epsilon=0$ is identical to the equations with no variation about the optimal impulse response then the extrema are described by equation 21. Then, using

$$\frac{\partial P(\epsilon)}{\partial \epsilon} = 0 = P_1 \tag{21}$$

the definition of S_0 we have

$$P_1 = \int_{-\infty}^{T}\delta h(T-\tau)\left\{\int_{-\infty}^{T} R_{\eta\eta}(\tau,\sigma)h(T-\sigma)d\sigma - \alpha S_0(T)s(\tau)\right\}d\tau = 0. \tag{22}$$

In this case, $\delta h(t)$ is arbitrary and so the bracketed term, $\{\}$, must be equal to 0. Then we have

$$\int_{-\infty}^{T} R_{\eta\eta}(\tau,\sigma)h(T-\sigma)d\sigma = \alpha S_0(T)s(\tau). \tag{23}$$

Notice that $\alpha S_0(T)$ is a constant and can be ignored as a scale factor. Also the autocorrelation of white noise is given by

$$R_{\eta\eta}(\tau,\sigma) = \frac{N_0}{2}\delta(\tau-\sigma) \tag{24}$$

where δ is a dirac delta function. Therefore, the integral drops out and we are left with

$$\frac{N_0}{2}h(T-\tau) = s(\tau). \tag{25}$$

$N_0/2$ is a constant and can again be ignored as a scale factor. With the change of variables, $T-\tau = t$, we have $h(t) = s(T-t)$. For the complex case we have $h(t) = s^*(T-t)$.

A.2 The Trigger Address Uncertainty

Several attempts were made to overcome the trigger address uncertainty. First the sampling boards were sent back to the manufacturer to attempt a hardware modification, but when they returned, the problem was not completely eliminated. There were two other options in using these boards. One idea takes advantage of the fact that some trigger address combinations had a higher probability of resulting in bad data than other addresses. Therefore, the trigger address could be checked and if the data has a high probability of being corrupted, the data is resampled. The process of checking the data and resampling as needed added 3% to the processing time. This doesn't seem like very much, but this arbitrarily changes the chirp rate of the data and so correlation is much less successful. Therefore the second option is used which is to set the data to zero if it was bad. This technique makes some range strips black in the image before azimuth compression, but this is just additional noise in the azimuth compression and the results are good.

References

Adams, R. E. W. et al. (1981). "Radar Mapping, Archeology, and Ancient Maya Land Use," *Science*, 25 September 1981, vol 213 pp. 1457-1463.

Arnold, David V. (1987). *Vector Quantization of Synthetic Array Radar Data*, Master's thesis, Brigham Young University.

Balanis, Constantine A. (1989). *Advanced Engineering Electromagnetics*, Wiley, New York.

Bruning, Schmidt, Alpers (1994). "Estimation of the Ocean Wave - Radar Modulation Transfer Function from Synthetic Aperture Radar Imagery," *Journal of Geophysical Research*, vol. 99 No C5 pp 9803-9815.

Curlander, John C. et al. (1991). *Synthetic Aperture Radar: Systems and Signal Processing*, Wiley, New York.

Drinkwater, Long and Early (1994). "Ers-1 Investigation of Southern Ocean Sea Ice Geophysics Using Combined Scatterometer and SAR Images," *IGARSS*.

Elachi, Charles et al. (1984). "Spaceborne Radar Subsurface Imaging in Hyperarid Regions," *IEEE Transactions on Geoscience and Remote Sensing*, **GE-22**, pp 383-388.

Farnett, Howard and Stevens (1970). "Pulse Compression Radar," *Radar Handbook*, chp 20, McGraw-Hill, New York.

Fitch, J. Patrick (1988). *Synthetic Aperture Radar*, Springer-Verlag, New York.

Harris, F. J. (1978). "On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform," *Proc. IEEE*, 66(1) pp 51-83.

Holcomb, Derrold W. (1992). "Shuttle Imaging Radar and Archaeological Survey in China's Taklamakan Desert," *Journal of Field Archaeology*, vol. 19, pp 129-137.

Klauder, Price, Darlington and Albersheim (1960). "The Theory and Design of Chirp Radars," *Bell Sys. Tech. Journal.*, vol 39, pp. 745-808, July 1960.

Levanon, Nadav (1982). *Radar Principles*, Wiley Canada.

Oppenheim, Alan V. and Willsky, Alan S. (1983). *Signals and Systems*, Prentice-Hall, New Jersey.

Nguyen et. al (1990). *SAR Data Catalog System User's Guide*, JPL unpublished.

Scollar, Irwin et al. (1990). *Archaeological Prospecting and Remote Sensing*, Cambridge University Press, Cambridge England.

Sherwin, Ruina and Rawcliffe (1962). "Some Early Developments in Synthetic Aperture Radar Systems," ***IRE Trans Mil. Electron***.,vol MIL-6, pp 111-115, April 1962.

Ulaby, Moore and Fung (1982). *Microwave Remote Sensing - Active and Passive Volume II Radar Remote Sensing and Surface Scattering and Emission Theory*, Artceh House, Norwood, MA.

Wehner, Donald R. (1987). *High Resolution Radar*, Artech House, Norwood, MA.

BYU SAR
A Low Cost, Compact Synthetic Aperture Radar

Bryan R. Jarrett

Department of Electrical and Computer Engineering

M.S. Degree, December 1995

ABSTRACT

A low cost, portable Synthetic Aperture Radar (SAR) was constructed that makes SAR available for a broader range of research fields. The system cost about $70,000, a hundredth of the cost of contemporary private SARs. The SAR demonstrated a range resolution of 2.5 meters with good sidelobe suppression and an azimuth resolution of 0.45 meters. The transmission power of the system was significantly lower than the other private SARs, but is sufficient for many applications. The resolution was demonstrated by imaging a field in Provo, Utah. The site contained a house, a fence, ground relief and other bright targets all of which showed up well in the image demonstrating the SAR's mapping ability. The Brigham Young University (BYU) SAR was deployed as an X-band SAR, but it could easily be modified (by changing antennas) to operate at Ku, C or L-band. The deployment took place from the bed of a truck, but the BYU SAR system is compact and light enough to be deployed from a variety of platforms such as a Cessna size airplane. By overcoming the cost and resolution barriers of previous SAR systems, the BYU SAR makes SAR technology accessible to the field of archaeology.

COMMITTEE APPROVAL:

David V. Arnold, Committee Chairman

David G. Long, Committee Member

Wynn C. Stirling, Graduate Coordinator