Time Domain Antenna Pattern Measurements

by

Andrew M. Predoehl

Thesis submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of
Master of Science
in
Electrical Engineering

APPROVED:

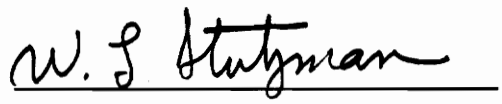

Dr. W. L. Stutzman, Advisor

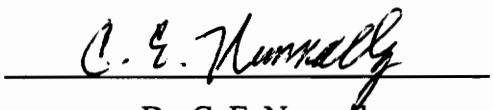

Dr. C. E. Nunnally

Dr. T. Pratt

June 1996

Blacksburg, Virginia

Key words: Antennas, Digital Signal Processing, Gating, Pattern, Time Domain

c.2

Time Domain Antenna Pattern Measurements

by

Andrew M. Predoehl

Abstract

Multipath on far-field antenna ranges causes distortion of antenna pattern measurements. The multipath components have differing path lengths and hence can be separated by illuminating the antenna under test with short-duration pulses. Alternatively, antenna measurements can be made in the frequency domain, and the Fourier transform can be used to relate the frequency-domain measurements to the antenna's time-domain response. The interference then can be removed with a time-domain gate, and transformed back into the frequency domain to yield improved CW antenna patterns. Virginia Tech has recently completed a major upgrade of their far-field antenna range and implemented a system to perform this data collection and data processing. This thesis describes the principles and implementation of the time-domain processing part of the system. Further, it demonstrates the validity of the method by showing the improvements in pattern measurement that have been achieved with the new system.

Acknowledgements

I want to thank several of the faculty, who have been very helpful to me and generous with their knowledge, time, and advice: Dr. William Davis, Dr. Charles Nunnally, Dr. Timothy Pratt, Dr. Jeffrey Reed, Dr. Brian Woerner, and especially Dr. Warren Stutzman, who has been a respectful employer and a supportive advisor, and who repeatedly hired me for projects that any human resources consultant would say I was unqualified to touch.

I am grateful to the Bradley Department of Electrical Engineering for the several teaching assistantships I have received from them, to the Satellite Communications Group and Center for Wireless Telecommunications for the research assistantships I have worked on, and to DuPont for the graduate fellowship given me.

I appreciate the help of the many students and staff in the Satellite Communications Group, Antenna Group, and EE department whom I've known over the years, especially that of Loretta Estes and Randall Nealy.

I am grateful to the many individuals who have made my graduate career (so far) as educational as it has been; I am even more grateful to those who have helped me cope with all the education I have received. It is my good fortune that they are so numerous that I cannot name them all here.

TABLE OF CONTENTS

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

The purpose of this chapter is to introduce many of the terms and ideas that are discussed in detail in this thesis. The discussion begins with a brief overview of antenna patterns and their importance, to justify the continuing interest in measuring antenna patterns. Then the focus shifts to the interference problem that arises in many measurement facilities, and the mitigation technique that is the thrust of this thesis.

1.1 Antenna Patterns And Interference

The antenna is a critical component in radio systems. It is a transducer that bridges the gap between electromagnetic (EM) waves and electrical circuits. See Fig. 1-1. Antenna theory and design continues to be an active research area. The recent surge in wireless communication has sparked interest in several different areas of antenna technology, from new, small antenna designs, to adaptive antennas for cellular base stations.

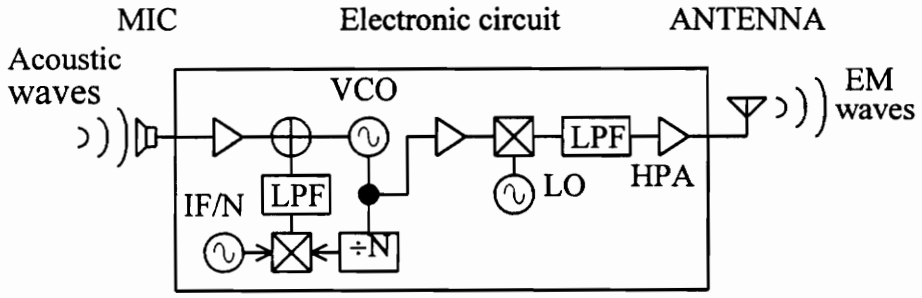

Figure 1-1: Diagram of a communications transmitter. Just as the microphone is a transducer of acoustic waves to electrical signals, so the antenna is a transducer from electrical signals to EM waves. (FM modulator circuit from [Cou90].)

The development of a new antenna usually involves fabrication of one or more prototypes, since many aspects of antenna behavior cannot be predicted with complete confidence, even with computer simulations. A prototype antenna can be tested in different ways, but often the most important test is that of measuring its *pattern*. (Terms in italics are explained in Appendix A.) The formal definition of antenna pattern can be found in [Stu81]; in essence, it is the antenna's power-transducing ability with respect to physical orientation. This idea will be further developed in Chap. 3.

In antenna design, pattern is often of critical importance. The relationship parallels that of the frequency characteristic of frequency-selective filter networks: although a filter can be evaluated in terms of its insertion loss, impedance, or group delay, usually the first thing a designer considers is its frequency characteristic. Likewise, there are many parameters that characterize an antenna (e.g., efficiency, impedance, and bandwidth), but pattern is almost always a consideration. Two examples illustrate this fact. The first is from the field of satellite communications, and the second pertains to radar.

Antennas play a very important role in satellite communications, both on spacecraft as well as on earth. As explained in [Pra86], geostationary satellites are often positioned only 2° apart along the geostationary orbit. Thus an earth station antenna must focus its transmitted power into a narrow beam, pointed directly at the target, and the antenna must be designed to keep the *illumination* of the neighboring satellites below a specified level. This level is specified (by the FCC) through antenna pattern. An earth station antenna pattern must stay within a prescribed "*sidelobe* envelope." Violation can result in large fines or even criminal conviction. Thus the engineer designing antennas for these applications will be very concerned with antenna pattern, especially sidelobe levels.

Many systems, such as radars, use phased array antennas. A phased array is an antenna made up of many smaller antennas, "elements," carefully arranged in a grid. The pattern can be "steered" electronically by establishing appropriate phase relationships between the elements. Unfortunately, certain factors such as fabrication on a substrate and

close proximity of the elements can combine to cause undesirable EM interactions between the elements. These interactions can have a profound effect on the array performance, causing "blind spots." The array cannot receive or radiate energy in the direction of the blind spot. This is a significant problem--a bad blind spot can render a phased array useless. Therefore blind spots must be detected as early as possible in the design. Often this is done by building a special prototype of the array and measuring its pattern, called the "active element pattern," which indicates the blind spot positions [Poz84]. Since the blind spots are hard to predict by other means, pattern measurement becomes an important part of the design of the phased array antenna.

The foregoing examples serve to illustrate the importance of pattern measurement. Antenna patterns can be measured with a number of techniques, which will be reviewed in Chap. 2. This thesis will focus on perhaps the most straightforward technique: measurement on a *far-field range*. This is performed by measuring the received power as the antenna is rotated about one axis while it is illuminated by a distant source. The resulting power level (usually expressed in relative or absolute dB) versus orientation forms the radiation pattern. (If the antenna is turned in only one axis, then strictly speaking only a subset of the pattern is measured, although we will not emphasize this distinction. The single-axis result is often called a *cut*.)

Ideally the *antenna under test*, or *AUT*, is excited with a plane wave, usually at a single frequency and polarization. In practice the excitation is generated by a *source antenna* located far enough from the AUT that the incident energy approximates a plane wave. Unfortunately, the source antenna tends to illuminate not only the AUT, but other objects as well. For example, if the range is located on a flat surface, the source antenna illuminates the ground. Some of the ground illumination will reflect into the AUT. Fig. 1-2 illustrates this idea (with some oversimplification). The result is that the measured "pattern" becomes the superposition of the true pattern and the pattern at a different orientation. In Fig. 1-2, the figure-eight shape symbolizes the true pattern of the AUT, excited by line-of-sight ray R1. The peanut shape beneath it represents the contribution of the

reflection ray R2, which excites the AUT from a different angle, generating a so-called "conical cut" component, which is in general different from the desired pattern. In a systems sense, most antennas are very linear, so superposition applies, and the two components add. The net effect is to distort the measurement, as shown on the right part of Fig. 1-2. Since antennas behave linearly, the network between the source antenna terminals and the AUT terminals will be referred to as the *test channel.*

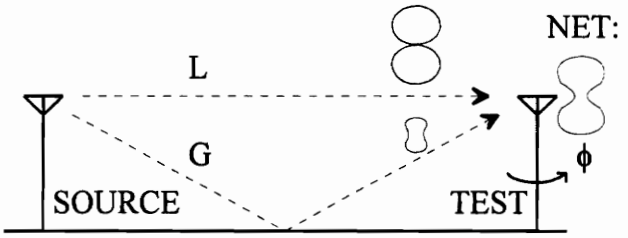

Figure 1-2: Illustration of how reflections introduce pattern measurement distortion. (Icons represent antenna pattern cuts--polar plots in the azimuth domain)

A practical far-field range can have many other sources of reflections besides the ground. For example, the Virginia Tech rooftop range has many potentially reflective objects within the source antenna's view. Fig. 1-3 is a photograph of the Virginia Tech far-field range, located on the roof of Whittemore Hall. There are many potentially reflective objects on this range, such as the roof of the building, the wall of the penthouse, the metal box visible in the foreground on the left, and the wall around the edge of the building. At higher frequencies these items tend to be less of a problem. However, below 3 GHz, the interference becomes significant. Because reflections can come from so many sources, it becomes impossible to account for all of them.

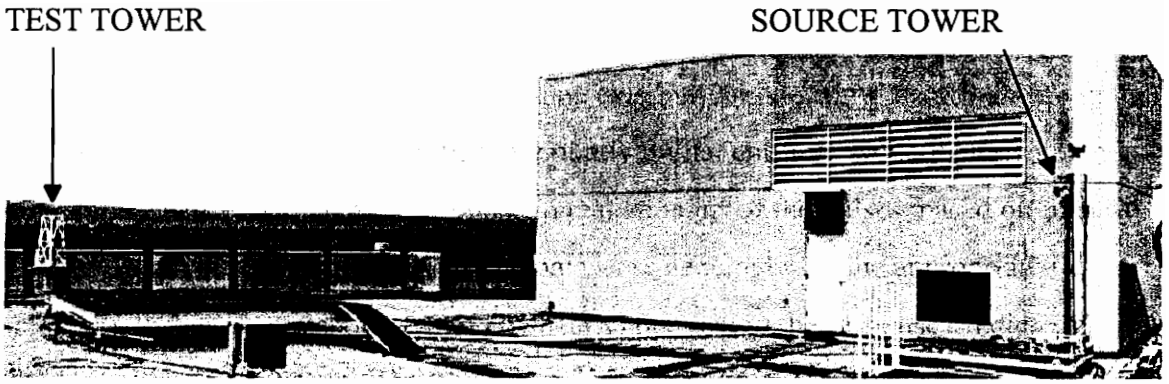

Figure 1-3: The Virginia Tech rooftop antenna range. The large central building is the Whittemore penthouse. Tower separation is about 20 m.

1.2 Time-Domain Gating As A Mitigation Technique

The preceding discussion assumes use of a single frequency--that is, an unmodulated carrier--also known as *continuous-wave (CW) excitation.* However, if the excitation is amplitude-modulated to shape the energy into a time-limited pulse, then the reflection components can be distinguished from the line-of-sight component, since the reflection paths are longer and hence experience more delay than the direct path. (Each meter of additional path length causes a 3.3 ns increase in the delay.) Fig. 1-4 illustrates this idea. Ray R1 is 18 m long, and a TEM wave requires 59.4 ns to go from the source antenna to the AUT along this path; whereas ray R2 is 1 m longer, and a TEM wave takes 3.3 ns longer to reach the AUT along this path. If the illumination is amplitude-modulated into short pulses, then the two components will be distinct. Fig. 1-5 shows what that might look like. A single pulse is transmitted from the source antenna, illuminating both the AUT and the ground. (The shape over the source antenna symbolizes the transmitted signal in the time domain: it is a sinusoidal carrier with amplitude modulation confining the energy into a short pulse.) Just as before, the energy travels to the AUT along a line-of-sight path and a reflection path. The times of arrival differ, though, because of the difference in path length. The symbols at each of the rays illustrate this: the pulse traveling

along the line-of-sight path arrives earlier than the pulse traveling along the reflection path. Thus the signal at the AUT (labeled RAW), shows two distinct pulses. We will refer to this signal as the *pulsed response* of the channel. The undesired component can be removed by multiplying the signal with a *time gate*. The result (labeled PROCESSED) is called the *gated response*. This operation is known as *gating*. To measure an entire antenna pattern using gating, the AUT must be rotated (continuously or in discrete steps) while this process is repeated.

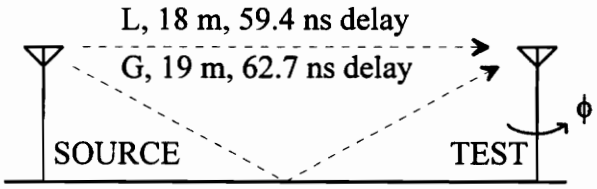

Figure 1-4: Dominant rays on the Virginia Tech antenna range. The reflection path has 3.3 ns more delay than the direct path.

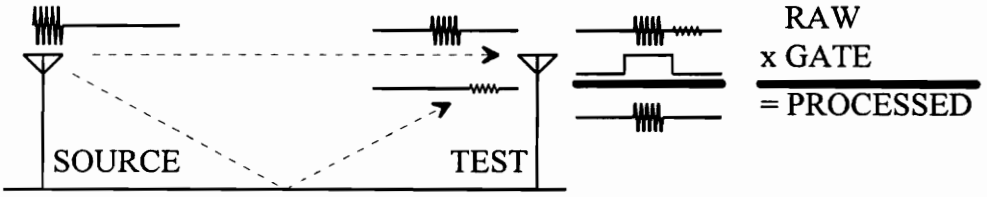

Figure 1-5: Illustration of time-domain gating.

Time gating can be performed using appropriate pulse transmitters and receivers. Transmitters with the pulse modulation and receivers with an adjustable time gate are commercially available. We will refer to this implementation as *hardware gating*, but it is not the focus of this thesis. For many reasons, it is desirable to use an alternate implementation, which we will call *software gating*. Software gating requires only a CW transmitter and receiver, and uses digital signal processing to determine the pulsed channel response. This thesis discusses the theory, design, development, and validation of the software gating system implemented at Virginia Tech.

A software gating system relies on the linearity of the channel to determine the channel's response to *synthetic pulses*, that is, hypothetical pulses that never exist as EM waves. The pulses are synthetic because they exist only in the computer; on the computer display, they look like pulses, but they are purely digital entities. Nevertheless, it is possible to determine the channel's response to real versions of the pulses, based on linearity. That is, if one of these synthetic pulse signals were generated via the appropriate RF hardware, the channel response would be nearly identical to the response predicted by the software gating system. The necessary assumptions are that the channel is linear and that the channel measurements are of high quality.

The information required to determine the pulsed response is a set of CW gain measurements (both magnitude and phase) made on the channel, across a range of frequencies, and repeated as the AUT is rotated in discrete steps. With these data the pulsed response can be determined numerically. Fig. 1-6 illustrates the computation of the channel's pulsed response. Synthetic pulses x(t) are decomposed into discrete frequencies, X(f), in a Fourier series. Channel measurements at the same discrete frequencies, T(f), characterize (at a fixed azimuth) the effects of the antennas, range reflections, and so on. The product of these frequency domain data yields the channel response to the synthetic pulse input, Y(f), represented in the frequency domain. The inverse Fourier series then relates the frequency-domain response to the time-domain response, y(t), which reveals the presence of interfering signals; the illustration suggests a line-of-sight component and a slightly-attenuated interference component. These concepts will be explained in greater detail in Chap. 3.

So far, the gating method of Fig. 1-5 has not been discussed--only the procedure to determine the pulsed response of the channel without having to transmit actual pulses. Fig. 1-7 illustrates the synthetic-pulse method of gating. A gate, g(t) is applied to the pulsed response y(t). The user chooses a gate function that passes the line-of-sight component and excludes the interference. This gate function is multiplied with the pulsed response in the time domain. Ideally, the product, the gated response, z(t), is the

interference-free response of the AUT (free of the reflection sources within the antenna range). The time-domain gated response then is decomposed into a Fourier series Z(f) that (ideally) indicates the AUT response over a range of discrete frequencies. If this procedure is repeated for each azimuth position of the AUT, then the discrete-frequency data in the gated response can be assembled into interference-free CW pattern cuts. Obviously a large number of cuts are generated at the same time--one for each frequency step.

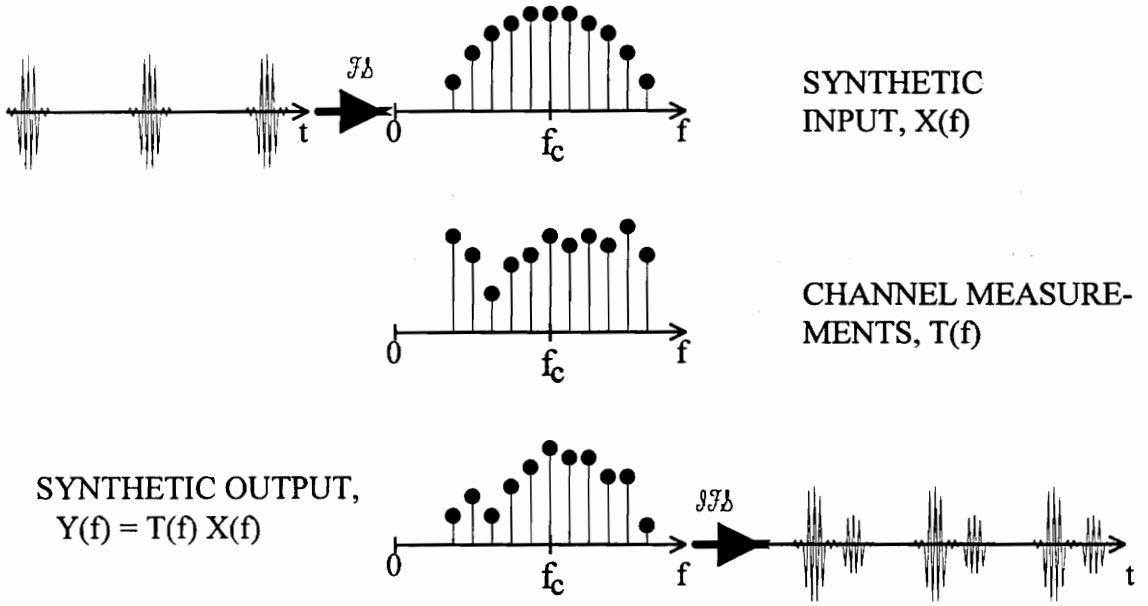

Figure 1-6: Illustration of the computation of the channel output due to a synthetic input. $\mathcal{FS}$ = Fourier series; $\mathcal{IFS}$ = inverse Fourier series

Three points need to be briefly addressed. First, it is perhaps not obvious that the spectrum of the gated response differs from the spectrum of the synthetic pulses; this will be established in Chap. 3. Second, Figs. 1-6 and 1-7 suggest that the relatively tapered spectrum of the synthetic pulses causes the gated response spectrum to be tapered. That is true here, but it can be circumvented easily, as will be seen in Chap. 3. Third, note that this method is "blind" with respect to azimuth; as a result, the same procedure must be repeated as the AUT azimuth angle is varied (usually in discrete steps). This is desirable if

the user intends to apply certain other error-correction techniques to the pattern data, such as pattern averaging.

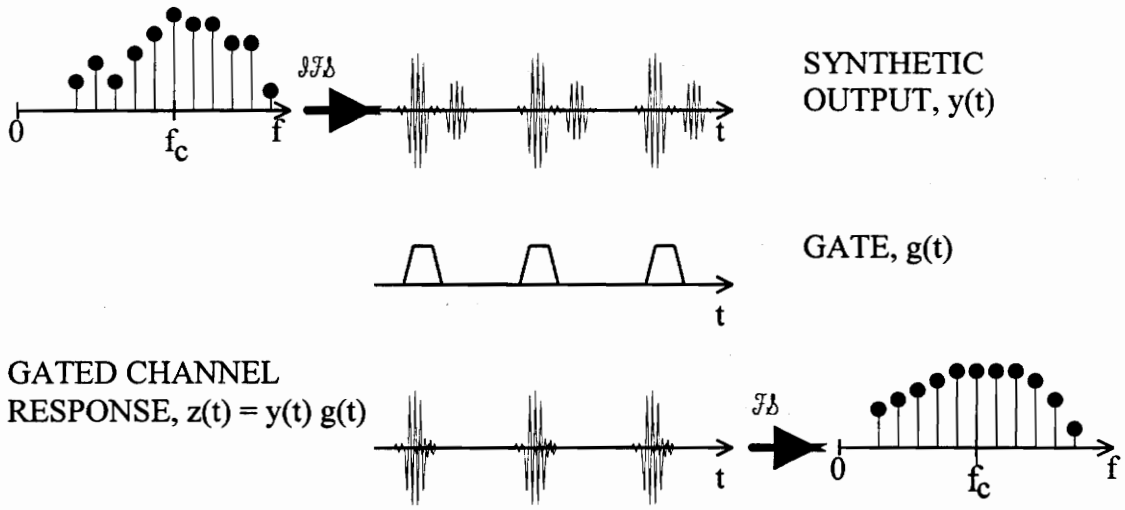

Figure 1-7: Illustration of the computation of the gated channel response.

As Fig. 1-6 and 1-7 indicate, the software gating system is conceptually more complex than the hardware gating system. But, the complexity is offset by several advantages. The most important is that is uses a CW transmitter and a CW receiver (capable of measuring both magnitude and phase). This is superior to a hardware gating system in a number of ways:

- CW transmitter and receiver technology is mature, accurate, and relatively easy to use.
- The transmitter can use uniformly high power levels at each frequency, and the receiver can use very narrow filters, yielding channel measurements with high signal-to-noise ratio.
- The Virginia Tech Antenna Laboratory, which sponsored the development, already owns the necessary equipment.

Other advantages include:

- Gated results can be related back to interference-free CW patterns via a Fourier series.
- Long after an experimental setup is disassembled, the AUT data can be re-loaded into the computer and gating parameters (such as gate position and width) can be changed.
- Pulsed response data in the time domain (without the gating) provide insight into the sources of range interference.

The main disadvantage of the software gating system is that the necessary data collection step is slow. A typical measurement might use 401 frequency steps at each of 360 azimuth positions, totaling 72360 data points. If the receiver uses a 27 Hz IF filter, which has a settling time of about 0.04 s [Ser84], then the data collection requires well over an hour. This is inconvenient for the user, of course, but it can cause a more insidious error: the measurement system can suffer gain and phase errors due to thermal drift. Other advantages and disadvantages will be discussed in following chapters.

There are many far-field ranges in use, most of which do not use time-domain techniques. These ranges often suffer reflection problems at certain frequencies. These ranges can be converted to serve at previously-unusable frequencies at a cost much less than that of designing and building a new antenna range. A software gating system requires a relatively modest investment: a PC and the proper software, plus an upgrade to a vector measurement receiver if the range is equipped with a scalar receiver.

To summarize, this thesis discusses a software gating system designed and implemented for the Virginia Tech rooftop range. The software gating system is an implementation of an interference-removal technique that relies upon Fourier series to relate frequency domain and time domain data at each azimuth angle.

Chapter 2 is a review of background information and the literature relevant to the methods discussed here. Although transformation between the time and frequency domains is not new, its application to antenna measurements is not widespread. The roots

of the technique lie in the field of radar research, and were popularized by those searching for novel network analyzer applications. The techniques have been applied to measurement systems more sophisticated than far-field ranges. Future directions for the concepts point towards higher resolution spectral estimation methods and interference cancellation techniques borrowed from spread spectrum communications. Other ideas for future work are found in Chap. 6. Chapter 3 presents the underlying transformation theory, building on the material in this chapter. Chapter 4 discusses *Range Runner*, the software package that was developed to run the antenna range and perform the time-domain gating in software. Chapter 5 discusses the validation of the new measurement system, and presents some sample data. The most notable effect of the gating system is that null depth has deepened by about 10 dB and the squinting common to many patterns has decreased substantially. Chapter 6 summarizes the contributions and presents ideas for future work.

CHAPTER 2

A REVIEW OF PATTERN MEASUREMENT TECHNIQUES

2.1 Introduction

Modern pattern measurement techniques are surveyed in this chapter; the reader is referred to Evans [Eva90a] for a detailed treatment. Three of the most popular facilities to measure antenna pattern are the outdoor far-field range (discussed in Chap. 1), the *compact range*, and the *near-field range*.

The outdoor far-field range is still a popular way to measure antenna patterns. The distance between the source and test antennas is ideally greater than $2D^2/\lambda$, where D is the maximum diameter of the test antenna, and λ is the free-space wavelength of the CW excitation. This criterion is based on the spherical phase error of the incident wave, assuming a point source. (Techniques have been developed to reduce this requirement, as in [Eva85], but they are beyond the scope of discussion.) The source and test antennas are typically elevated on towers, and at least one positioner is used to control the azimuth of the test antenna. Ground-bounce interference is sometimes minimized by using fences (to scatter the energy away from the test antenna) or microwave absorber (to convert the energy into heat). These techniques can cause improvement in some cases, but as discussed in Chap. 1, if there are many reflective objects on the range, the ground-bounce interference might not be the only source of difficulty. These conventional techniques become impractical when one cannot determine the origins of the reflections.

A *compact range* is a measurement facility that uses a large reflector to produce plane-wave illumination. This reduces the real-estate expense of the range, at the cost of

a more complex source antenna. The source antenna (usually an offset-fed reflector) is designed to illuminate the test zone with a wave of uniform amplitude and phase close to the antenna aperture (rather than via a small source antenna located a long distance away). Again, at least one positioner is required to turn the test antenna. Compact ranges are indoor facilities, and thus the range designer must deal with reflections from six sides of a box. To deal with this problem, the room surfaces are covered with microwave absorber, and the feed and reflector are carefully designed to prevent interference components from contaminating the test antenna illumination. Techniques similar to the software gating process have been applied to compact ranges, and they will be discussed below; however, they are not the focus of this thesis.

A *near-field range* is an antenna measurement facility that employs a radically different approach to pattern measurement. Near-field ranges sample the EM field near the test antenna, and use mathematical transformations to determine the antenna's far-field pattern from the data. The technique is naturally resistant to reflection interference, and thus will not be discussed.

Evans [Eva90b] provides an overview of time-domain gating techniques. Of course, the use of gating is not restricted to far-field ranges; although gating can be applied to all three of these pattern measurement methods, often this technique is an appropriate solution to the interference problem on far-field ranges, because of the longer distances that are a part of the range geometry. Evans establishes the advantages of gating and the disadvantages of pulsed excitation, thus motivating the use of CW excitation to determine the channel's pulsed response.

2.2 Early Application Of Time-Domain Techniques To Antennas

The extraction of frequency-domain behavior from time-domain data is examined, both in theory and practice, by Nicolson [Nic68]. The application is the measurement of S11 and S12 for wideband microwave devices such as couplers and attenuators. A

broadband source such as a pulse or step generator produces a transient response (either transmission or reflection), the samples of which then are Fourier-transformed to yield frequency domain results. Lamensdorf and Susman [Lam71] apply the above technique to antenna measurements, in order to determine the wideband response of antennas. That is, on an ideal far-field range, this technique yields antenna response (gain or pattern) to wideband excitation, which reveals how the antenna will behave in a wideband system.

The first researchers to apply gating to antenna measurements were Burrell and Jamieson [Bur73]. They used pulse excitation and a sampling oscilloscope to make broadband pattern measurements. By sampling only within a finite time window, a time gate is implicitly established. Reduced-interference CW patterns then can be estimated using a Fourier transform on the time-domain measurements.

2.3 Time-Frequency Transformation In Radar Literature

Many of the issues involved in implementing a gating system with CW excitation are explored in the literature for radar, including the procedures for transforming back and forth between the frequency and time domains efficiently, the use of frequency-domain and time-domain tapers, and the limitations of practical measurement system. All the theoretical work necessary for software gating is found in this literature; the difference is the application.

Mensa [Men84] discusses the transformation between the time-domain and frequency-domain in the context of radar cross-section (RCS) measurements. He shows how the measured frequency domain response (whether FM or stepped-frequency CW excitation) can be transformed into the time domain with an FFT, how the interference effects can be removed with a tapered gate, and then how the time-domain results are transformed back into the frequency domain to better estimate the target's narrowband (CW) RCS. Dominek, Peters, and Burnside [Dom87] discuss how the EM scattering mechanisms for a radar target can be identified and separated from one another by using time-

domain techniques. Mechanisms such as specular reflection, edge diffraction, and creeping wave diffraction can be observed when the scattering is shown in the time domain, since the mechanisms tend to be spatially localized, and hence localized in the time domain. Using CW excitation to measure the "scattering spectrum" from various radar targets (e.g., metal spheres, disks), the measurements may be tapered and transformed into the time domain; then a single mechanism can be isolated with a tapered gate, and the results transforms back to the frequency domain to obtain the spectrum of the desired mechanism. These steps all parallel those used in the software gating system described in this thesis.

2.4 Network Analyzer Applications

The antenna range can be viewed as a channel consisting of the source antenna, the range, and the AUT. It is a two-port network (albeit a physically large one), with a set of scattering parameters, and thus may be characterized with a network analyzer. Measurement of S21 as azimuth varies represents that pattern of the AUT. The HP 8510 network analyzer is well-known for its time-domain capabilities, including software gating, and has been the subject of several application-oriented antenna measurement papers, although we will discuss only one. Boyles [Boy85] is perhaps the first to discuss the application of the gating capabilities of the HP 8510 network analyzer to antenna pattern measurement. He presents an example using a dipole at 115 MHz. S21 was measured with CW excitation from 50 MHz to 300 MHz in 1.25 MHz steps. The HP 8510 uses Fourier transforms to determine the channel's pulsed response, and Boyles shows both the frequency-domain and time-domain response of the antenna. Since the antenna is relatively narrowband (thus the channel is narrowband), the time-domain response shows overlap between the direct energy and the ground-reflection interference (determined by inspection). He uses a tapered time gate to remove most of the ground reflection

response, yielding a dipole pattern showing nulls down to -30 dB, and agreeing with the theoretically-predicted pattern much better than did the ungated measurements.

2.5 Errors Due To Time-Gating

Fig. 1-5 is an oversimplified representation of an antenna's response. Practical antennas have limited bandwidths and therefore have potentially complicated and lengthy, even infinite, impulse- and step-responses; however, a time gate is only of finite duration. Thus gating rejects interference, but it also truncates the AUT's time response. This introduces error into the measured pattern. Henderson, James, Newham, and Morris [Hen89] have examined the errors that occur due to the gating process by modeling the time-domain response of wideband antennas with lumped-parameter bandpass filters, and then computing (analytically or numerically) the impulse response of the models. The gating error is caused by the truncated portion of the impulse response, that is, the part of the impulse response that falls outside the gate. As would be expected, the lower cutoff frequency of the filter strongly affects the duration of the impulse response and hence the amount of gating error for a given gate size. Given the geometry of an antenna range, one can estimate the minimum delay between the line-of-sight and interference responses, and hence the maximum time gate that rejects interference. Along with this information, if one has the approximate frequency response of an AUT, one could model the antenna with a bandpass filter; then the gate of maximum duration can be applied to the impulse response of the model, yielding an estimate of the error that will occur in the gating process.

2.6 Compact Range Applications

A compact range has several potential sources of nearby reflections, which can distort pattern measurements. The inside of a compact range is usually covered with a layer of microwave-absorbing foam, which can reduce reflections by 20 to 40 dB

[Eva90a] but not completely eliminate them. Other techniques are used to reduce interference, such as careful feed design to keep the feed's backlobes low. Gating, too, can be used in a compact range to combat reflections and improve the pattern measurement accuracy. Hess et al. [Hes88] experimented with pulsed excitation and hardware gating to improve the test zone illumination in a compact range, but found that the gating provides no significant improvement over the use of more traditional techniques. When interference is deliberately introduced into the range (e.g., the source feed is modified so that its backlobes are increased, and some of the absorber is removed), then they found that gating improves the range performance to a level similar to that of the unimpaired range. Thus, they conclude that for a well-designed compact range, the additional complexity of gating may be superfluous.

Both hardware gating and Fourier-transform techniques can be employed simulateously, as discussed in [Lam87]. The application is the measurement of the far-field pattern of two large (8' dia.) reflector antennas with the Ohio State University compact range. Pulsed excitation and hardware gating are used to eliminate the effects of wall reflections in the compact range. In contrast to [Hes88], the authors note that hardware gating is an important feature of the compact range, and is partially responsible for the high dynamic range they claim for their measurement system. The authors use the pulses as if they were CW excitation; that is, they step the pulse's center frequency across a band, and collect data that approximates the channel's CW response. (If the pulses are relatively long, the approximation will be good.)

The pseudo-CW response is transformed into the time domain, and the results show the mechanisms at work inside the AUT: contributions from subreflector spillover, single-bounce energy, double-bounce energy, etc. The point here is not that they use software gating to exclude these effects, but rather that the information shown in the time-domain response permits them to evaluate how the antenna is working. This demonstrates that interference reduction is only one application of time/frequency transformations in antenna measurements.

2.7 Future Work

Research in interference detection and cancellation continues, comprising techniques that go beyond the limitations implicit in Fourier-transform type approaches, and hence beyond the method examined in this thesis. One area of interest involves spread-spectrum interference cancellation techniques. The technique discussed in [Zeg95] rests on certain assumptions about the channel and the received signal; these will need to be tested in practice. Another area of activity is the application of "superresolution" techniques (modern spectral analysis), such as MUSIC, to frequency-domain data [Oga91]. These methods bypass the resolution limitations of Fourier-transform type approaches, and perhaps could be very effective on narrowband antennas. That is, they provide much higher resolution in the time domain than the Fourier transform does. However, it is unclear whether one can use the time-domain results to cancel the interference.

In conclusion, an antenna range is a facility for measuring antenna patterns. Most ranges actually or potentially have problems with range multipath, which interferes with pattern measurement. However, this interference can be reduced by gating. The necessary theoretical foundation for the software gating system presented here came from the literature of radar. It was adapted to network analyzers and soon applied to antenna pattern measurements, on both far-field and compact ranges. Gating can significantly improve antenna pattern measurements. Gating can cause errors, too, and in general the appropriate gate size cannot be prescribed without the use of one's judgment. Current work in interference reduction involves non-Fourier-transform type approaches.

CHAPTER 3

PRINCIPLES OF SOFTWARE GATING

This chapter presents the principles underlying the software gating system. Both signal processing theory and measurement issues are involved, and the latter is where the discussion begins.

3.1 Model And Measurement Issues

The experimental setup of Fig. 1-3 can be modeled as shown in Fig. 3-1. The test channel is regarded as a causal, linear two-port network; furthermore, it can easily be made adequately time-invariant by fixing the AUT azimuth angle, ϕ. Therefore it has scattering parameters [Bal81]. (In practice, ϕ is stepped slowly, but the energy storage between steps is negligible, so the network can be regarded time-invariant.) The scattering parameters are illustrated in Fig. 3-1 (b). The parameter of greatest interest is the transmission coefficient S21, which changes with azimuth ϕ and frequency f, and so is written as S21(ϕ, f). (Recall that all RF measurements use CW excitation, not pulses.)

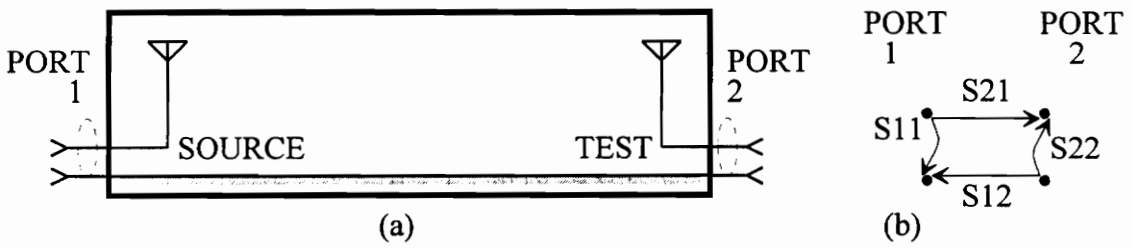

Figure 3-1: Test channel portrayed as (a) a two-port network (AUT azimuth fixed at ϕ), and (b) a flow-graph diagram of scattering parameters.

The typical means of measuring the transmission coefficient is shown in Fig. 3-2. A CW RF transmitter is coupled to port 1, and a CW vector receiver measures the amplitude and phase of the signal at port 2. Phase measurements are needed to perform the software gating procedure, hence the receiver needs a phase reference signal. This can be implemented in a number of ways. Fig. 3-2 shows the way used in the system considered in this thesis. A splitter, S1, is used to couple some of the signal feeding port 1 to serve as a reference signal. One disadvantage of this method is that an attenuator, A2, is required between the splitter and receiver to keep the reference signal level from saturating the receiver. About 40 - 60 dB of attenuation is usually necessary.

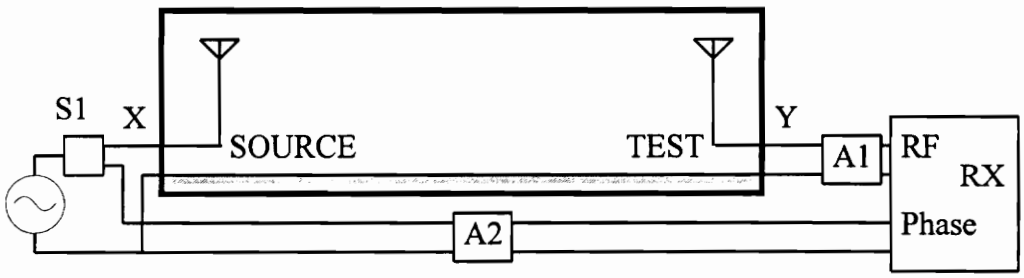

Figure 3-2: Measurement of transmission coefficient

Attenuator A1 is optional but beneficial; it prevents cable reflections from appearing in the time-domain signal. Attenuation of 10 - 20 dB should be adequate. As is well-known, if an EM wave encounters an impedance discontinuity, a reflection occurs. (This phenomenon causes the range multipath, too.) Since in general the impedance of the AUT differs from the impedance of the measurement system, and reflections occur in the cabling, it is prudent pad the line to suppress cable reflections that will obscure the multipath reflections. If the reflections due to changing antenna impedance can be ignored, then the channel output, Y, can be written in terms of S21 and channel input X:

$$Y(\phi, f) = S21(\phi, f)\, X(f).$$

Inevitably, there is a large path loss between the source and test antennas; only a small fraction of the power that enters port 1 is available at port 2. It is therefore difficult to make an exact measurement of S21. Fortunately, it is rarely necessary. Pattern measurements follow from the relative difference between $S21(\phi, f)$ coefficients as azimuth and frequency vary. This is measured by fixing the input signal level and measuring the change in output signal levels. For example, suppose $Y(\phi_1, f)$ is 1 dB higher than $Y(\phi_2, f)$, and 60 dB lower than input X(f). It is difficult to measure the 60 dB difference between X(f) and $Y(\phi_1, f)$, and usually not required; only the 1 dB difference between $Y(\phi_1, f)$ and $Y(\phi_2, f)$ is measured. Therefore in practice S21 is unknown in an absolute sense, and only the relative difference between $S21(\phi_1, f_2)$ and $S21(\phi_3, f_4)$ is stored. With the receiver used in the Virginia Tech system (Scientific-Atlanta model 1780), the receiver generates a power measurement referenced to an unknown, but fixed, absolute power level, and proportional to the output power. Even though the transmission coefficient is only partially known, we will nevertheless refer to these measurements as S21 measurements in this thesis.

The relationship between S21 and relative pattern is simple; that is, if the antenna pattern is normalized, and not referenced to *antenna gain*. Techniques for gain measurements are beyond the scope of this thesis; we will consider normalized patterns only. A normalized pattern indicates only the relative change in gain versus orientation, and therefore by convention the peak of the pattern is scaled to 1, or 0 dB. Given a set of S21 measurements at a particular frequency, $\{S21(\phi, f) : 0° \le \phi < 360°\}$, the normalized pattern is the set of measurements divided by the maximum value in the set; that is,

$$p_f(\phi, f) = S21(\phi, f) / \max_{\phi} \{S21(\phi, f) : 0 \le \phi < 360\}, \qquad (3.1)$$

where $p_f(\phi, f)$ denotes a normalized pattern cut at frequency f. Normalization introduces one difficulty. A full antenna pattern may be expressed as a function of spherical coordinates, (θ, ϕ). Some ranges are equipped to vary both θ and ϕ, although this thesis will

focus on single-axis cuts of the pattern. In this case, the user will not know if he has made a measurement in the (θ, ϕ) direction of peak gain. Nevertheless, each cut is normalized so the peak is scaled to 0 dB. Different cuts of the same pattern can be compared to each other if additional information is provided with each cut, namely, a *reference level*, the power measured at the receiver at the cut peak. Fig. 3-3 illustrates this idea. The surface represents a complete antenna pattern over all values of (θ, ϕ). Two cuts, A and B, are shown to the right. Cut A intersects the pattern peak, but cut B does not. If the reference level of cut A were, for example, $P_{ref, A}$ = -17 dBm, and the reference level of cut B were $P_{ref, B}$ = -29 dBm, then the AUT has more gain in the sidelobes of A than in the sidelobes of B, despite the appearance of the normalized patterns. Reference levels have additional utility. They are part of the procedure for making gain measurements, and they are useful for comparing cuts of differing frequencies just as they are for comparing cuts of different axes. One requirement for the software gating system is that it must preserve the accuracy of the reference level measured with each cut.

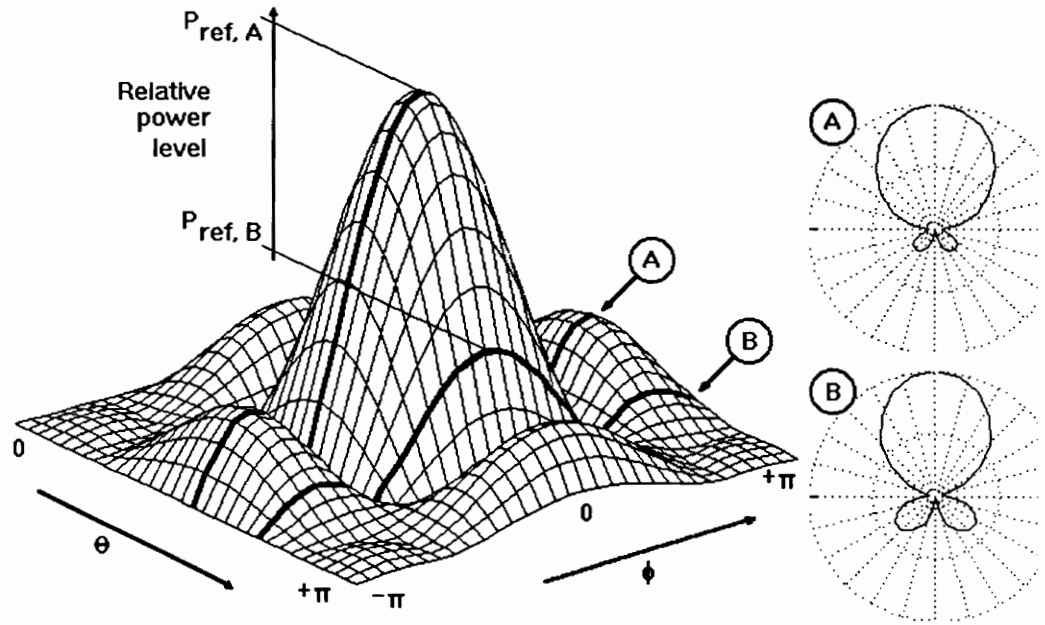

Figure 3-3: Illustration of pattern normalization. Cuts A and B are each normalized. The relationship between A and B is lost unless a reference level for each cut is retained.

Dynamic range is an important consideration for measurement systems. One factor that determines dynamic range is the system noise floor. Thermal noise, since it is proportional to bandwidth, can be reduced by narrowband filtering; therefore it is advantageous to filter the channel output while measuring S21. Unfortunately, narrowband filters have a long impulse response, the length of which is inversely proportional to the filter bandwidth. Of course, $S21(\phi, f)$ should indicate the channel response to CW input, and ideally, CW input has the form $\cos(2\pi ft)$, which extends back in time from $t=-\infty$. In practice the channel is excited with a sinusoid of only finite duration, which must be long enough to permit the transient response to die out. This transient period is very short for the antennas under consideration, and is dominated by the IF filter used in the measurement receiver. For example, the SA 1780 receiver used in the system under discussion gives the user a choice of filters, with 3 dB bandwidths ranging from 1600 Hz to 7 Hz; the transient period of these filters ranges from 9 ms to 525 ms [Ser84]. This modest delay, however, can lead to a significant total measurement duration.

One consequence of a long measurement duration is that the temperature of the system components can change, especially if the sun rises or sets during this period. The SA 1780 receiver hardware is indoors, except for measurements above 1 GHz, a harmonic mixer, which downconverts the antenna RF to a 45 MHz IF, is located outdoors, often exposed to sunlight. Furthermore, on the Virginia Tech rooftop range, over 20 meters of coaxial cable run between the harmonic mixer and the appropriate receiver port; temperature changes may cause the length of this cable to change, thus causing the apparent delay in the channel to vary with time. It will thus be desirable to make measurements in weather conditions that remain relatively constant during the duration of the experiment. Fortunately, the local oscillator (LO) unit of the receiver is kept indoors; its output feeds the harmonic mixer via the 20 m coaxial cable just described. The resulting thermal stability no doubt is beneficial for the leveling and phase stability of the LO signal.

3.2 Multipath Effects

Range multipath distorts CW antenna pattern measurements. In Fig. 1-2, the distortion effects are illustrated with a ray diagram. (A description of multipath effects that uses rays should be sufficiently accurate for a far-field range, since presumably the range is large enough to justify the geometric approach, and uncluttered enough to keep other propagation mechanisms insignificant.) The ray components will interfere constructively or destructively when they combine at the AUT, depending on the relative differences in phase. Thus a ray diagram predicts that the interference-corrupted pattern measurement consists of a line-of-sight pattern plus one or more components of unknown pattern (due to the different direction of arrival), enhancing and diminishing the desired pattern (depending, in part, on differences in path length.) One must also regard the polarization of each multipath ray as unknown, since reflection coefficients are polarization-dependant.

These components can be separated from one another with a set of CW S21 measurements across a wide bandwidth because each interference ray propagates along a path that is longer than the line-of-sight ray; hence each multipath contribution arrives with a delay distinct from that of the line-of-sight component, and each contribution has a distinct phase-vs.-frequency characteristic. A multipath component arrives at an angle much different from that of the direct ray and with altered polarization; the amplitude will be unknown, and perhaps frequency-dependent. Regardless of all this, the additional delay in the indirect path gives the multipath component a steeper phase-vs.-frequency curve than that of the direct ray. If this curve is substantially steeper, then the interference effect can be removed via a correlation-style integral. For example, let $S21(\phi, f)$ be a transmission coefficient measurement on a channel like the one portrayed in Fig. 1-2, having a significant line-of-sight component $L_1(\phi, f)$ and ground-bounce component $G_1(\phi, f)$. Since the channel is linear,

$$S21(\phi, f) = L_1(\phi, f) + G_1(\phi, f). \tag{3.2}$$

If we let τ_L and τ_G represent the gross delay in each path, proportional to the path length, then we can factor the consequent phase shift from L_1 and G_1 as follows:

$$S21(\phi, f) = L_0(\phi, f) \exp(-j2\pi f\tau_L) + G_0(\phi, f) \exp(-j2\pi f\tau_G), \tag{3.3}$$

where L_0 and G_0 are the rays' frequency responses less the phase shift due to path lengths. To isolate the line-of-sight antenna response, we multiply both sides by $\exp(j2\pi f\tau_L)$ and integrate:

$$\begin{aligned} s21(\phi, \tau_L) &= \int_{-\infty}^{\infty} S21(\phi, f)\ \exp(j2\pi f\tau_L)\ df && (3.4)\\ &= \int_{-\infty}^{\infty} L_0(\phi, f)\ \exp(-j2\pi f\tau_L)\ \exp(j2\pi f\tau_L) + G_0(\phi, f)\ \exp(-j2\pi f\tau_G)\ \exp(j2\pi f\tau_L)\ df \\ &= \int_{-\infty}^{\infty} L_0(\phi, f)\ df + \int_{-\infty}^{\infty} G_0(\phi, f)\ \exp[j2\pi f(\tau_L - \tau_G)]\ df && (3.5)\\ &= l_0(\phi, 0) + g_0(\phi, \tau_L - \tau_G). && (3.6) \end{aligned}$$

In practice both L_0 and G_0 are band-limited measurements, and hence the integrals cover only a finite frequency range. If L_0 and G_0 are relatively flat with respect to frequency, then the second term in (3.5) will oscillate around zero, whereas the first term will be significant. Thus we can separate the ground-bounce component G_0 from the line-of-sight component L_0, regardless of the pattern excited by the ground-bounce ray. We require that L_0 and G_0 be fairly flat in the frequency range of interest, so that the first term in (3.5) outweighs the second term; however, L_0 and G_0 do not need to be strictly constant to do so. The frequency range of the measurement must be relatively large compared to $(\tau_L - \tau_G)^{-1}$ so that the exponential in the second term of (3.5) is not a significant factor.

This argument can be extended to consider a large number of ray interference components. As long as each interfering ray has an adequately well-behaved frequency response, and the delay due to extra path length is significant compared to the measurement bandwidth, the interference can be separated in the same way.

Of course, (3.4) is the inverse Fourier transform (IFT) of S21. The transmission coefficient is expressed in the time domain as $s21(\phi, \tau_L)$. Eq. (3.6) expresses it as the

superposition of two responses, l_0 and g_0, which are the time-domain impulse responses of the two rays, minus the delay due to path length. Thus $l_0(\phi, 0)$ is the start of the time-domain response of the line-of-sight ray, although it will not ordinarily be the peak. The ground-bounce component is still present, although it is evaluated at time $\tau_L - \tau_G$. Note that $\tau_G > \tau_L$, so $\tau_L - \tau_G$ is negative. Thus the ground-bounce component, which is causal, ideally contributes nothing to the transmission coefficient evaluated at time τ_L.

In (3.2) - (3.6), ideally the measurement bandwidth of S21, hence that of L_0 and G_0, is infinite, and the time-domain responses s21, l_0, and g_0 are impulse responses. Of course, since the S21 measurements are band-limited, isolation is imperfect, and the line-of-sight component is contaminated by the ground-bounce component. This contamination can be explained in both the frequency domain and the time domain. In the frequency domain, if $\tau_L - \tau_G$ is small, the exponential in the second term of (3.5) will be significant, which will cause the ground-bounce response to affect the line-of-sight component. In the time domain, the effect is explained via the time-domain equivalent of spectral leakage; s21, and hence l_0 and g_0, are smeared in the time domain because they are excited (convolved) not with an impulse, but a band-limited approximation, which we shall call r(t). For example, if S21 were measured in a sweep of spectrum with bandwidth B, centered at frequency f_c, then r(t) is the IFT of the excitation signal, which is flat from f_c-B/2 to f_c+B/2:

$$r(t) = \int_{f_c-B/2}^{f_c+B/2} 1 \cdot \exp(j2\pi ft)\, df = \frac{\sin(\pi Bt)}{\pi t} \exp(j2\pi f_c t) \ . \qquad (3.7)$$

Thus r(t) is a pulse with a sinc envelope modulating a complex exponential. The sinc envelope, of course, peaks at time t=0, but nevertheless extends from negative to positive infinity. Thus the rays are excited by an input that is imperfectly time-localized. Although l_0 and g_0 are causal, the corresponding time responses extend from negative to positive infinity, and interfere with each other. But, under the appropriate conditions, the interference will be reduced. Thus we can summarize (3.4) - (3.7) as follows:

$$s21(\phi, t) = l_0(\phi, t) * r(t - \tau_L) + g_0(\phi, t) * r(t - \tau_G) + \ldots \quad (3.8)$$

where the ellipses indicate where other multipath responses can be added to the model. The asterisks denote linear convolution. In the limiting case of infinite measurement bandwidth, r(t) reduces to a delta function.

The goal of this discussion is to introduce how the multipath interference causes complicated distortion of the desired pattern measurement, but, given a wideband S21 measurement of the channel, the differing phase-versus-frequency behavior of the interference components may permit the line-of-sight component to be isolated--at least partially. The procedure for isolating the line-of-sight component is more complex than that introduced here, but the fundamental limitations that arise from (3.5) in the frequency domain, or (3.8) in the time domain, ultimately constrain how much improvement is possible from a Fourier-transform solution.

3.3 Wideband Characterization

S21 measurements across a wide bandwidth are achieved with narrowband measurements at uniformly spaced discrete frequencies. Although it is possible to make fairly accurate transmission coefficient measurements across a continuous bandwidth by using appropriate swept-frequency equipment, this approach was not considered because it would require extra equipment, it would make the gating software much more complicated, and it would not confer any significant advantages. Thus the approach taken in the software gating system is to characterize the channel at uniform frequency steps. The frequency limits, bandwidth, frequency step, and number of steps are all interrelated parameters:

$$B = f_U - f_L \quad (3.9)$$

$$\Delta f = B / (N-1) \quad (3.10)$$

$$N \in \{\text{odd integers} > 1\} \quad (3.11)$$

$$N = 1 + B/\Delta f \tag{3.12}$$

$$f_U = f_L + (N\text{-}1)\,\Delta f \tag{3.13}$$

where B is the bandwidth of the measurement, f_U and f_L are the upper and lower frequency limits of the measurement, Δf is the frequency step size, and N is the number of frequency steps, which must be an odd integer larger than one. N is odd for reasons of convenience explained later. This procedure is repeated at discrete azimuth steps, typically starting at ϕ=0°, incrementing by $\Delta\phi$ at the end of each frequency sweep, up to M total azimuth positions. Thus, a typical full set of measurements consists of a measurement first of S21(0°, f_L), then of S21(0°, f_L+Δf), then S21(0°, f_L+2Δf), and so on, up to S21(0°, f_L+ (N-1) Δf), which is the same as S21(0°, f_U) by (3.13); then it continues with S21(0° + $\Delta\phi$, f_L), and so on. The set of measurements can be organized in a two-dimensional array **S**:

$$\mathbf{S} = \begin{pmatrix} S21(0, f_L) & S21(0, f_L + \Delta f) & S21(0, f_L + 2\Delta f) & \cdots & S21(0, f_U) \\ S21(\Delta\phi, f_L) & S21(\Delta\phi, f_L + \Delta f) & S21(\Delta\phi, f_L + 2\Delta f) & \cdots & S21(\Delta\phi, f_U) \\ S21(2\Delta\phi, f_L) & S21(2\Delta\phi, f_L + \Delta f) & S21(2\Delta\phi, f_L + 2\Delta f) & \cdots & S21(2\Delta\phi, f_U) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ S21((M-1)\Delta\phi, f_L) & \cdots & \cdots & \cdots & S21((M-1)\Delta\phi, f_U) \end{pmatrix} \tag{3.14}$$

We will refer to this 2D array either as **S** or $\{s_{i,j}\}$. **S** has M rows and N columns, and a typical element $s_{i,j}$ = S21((i-1) $\Delta\phi$, f_L + (j-1) Δf). (It would be misleading to call **S** a matrix, since it is never used in a matrix operations; it is just an array in which data are stored.)

As the channel excitation frequency increases from f_L to f_U, ideally only the ray interference and the characteristics of the AUT cause the amplitude and phase of the output to change. To implement this, the following conditions must be met:

- The CW source must be "leveled"

- The reference channel must be nondispersive (i.e., delay independent of frequency)
- The source antenna must be relatively flat in both amplitude and phase, and
- The receiver, including the mixers, local oscillators, detectors, and other components, must have flat amplitude and phase characteristics.

In the case of the SA 1780 receiver, the frequency range of operation is divided into fixed "bands" of operation (for example, 1-2 GHz, 2-4 GHz, etc.) over which the receiver is acceptably flat. Measurement across a band boundary requires a calibration to determine the relationship between the two relative power levels. For example, a sweep from 1.5 to 2.5 GHz crosses the 2 GHz band boundary. While keeping all other factors the same, one can, while the receiver uses the 1-2 GHz band, measure the channel output at 2 GHz, $Y_{1\text{-}2}(\phi, 2\ \text{GHz})$; and then again while the receiver uses the 2-4 GHz band, $Y_{2\text{-}4}(\phi, 2\ \text{GHz})$. Then the factor

$$k(\phi) = Y_{2\text{-}4}(\phi, 2\ \text{GHz}) / Y_{1\text{-}2}(\phi, 2\ \text{GHz})$$

can be multiplied with subsequent S21 measurements for frequencies above 2 GHz, which will reference them to the same level as the measurements below 2 GHz. Strictly speaking this calibration factor k should be independent of orientation ϕ, but in the system under consideration it is recomputed at each azimuth step. This keeps the code simple, it compensates the effect of thermal drift in the receiver, and it only slightly increases the measurement time. This can cause problems, too though: the channel must have output significantly above the noise floor at each of these band boundaries. If, for some reason, the channel output at 2 GHz is beneath the noise floor, then k will be meaningless and will cause problems--especially if the AUT has nulls that steer with frequency.

3.4 Frequency-To-Time Transformation

As introduced in Sec. 3.2, one approach to separating the line-of-sight component of S21 from the interference components is to employ a correlation-style integral that enhances the line-of-sight response and tends to cancel the interference responses. We saw that this procedure is the familiar inverse Fourier transform. In this section we will discuss these principles in more depth, and show how the S21 measurement serves as a basis for determining the channel output to a variety of inputs.

3.4.1 Synthetic Inputs And Outputs

Given a set of channel measurements **S** it is possible to compute the channel output to certain inputs--for example, a sinusoidal input at frequency f_L. See Fig. 3-2. If the channel input is of the form

$$x(t) = A\cos(2\pi f_L t+\theta),$$

then the channel output is

$$y(\phi, t) = |S21(\phi, f_L)|\, A\cos[2\pi f_L t + \theta + \angle S21(\phi, f_L)].$$

Admittedly, this is a trivial application of scattering parameters. In fact, it is trivial to determine the output corresponding to any sinusoidal input at frequency $f_L + k\,\Delta f$, $0 \le k < N$. Note that one may hypothesize a channel input based on two or more of these sinusoids, superimposed to form a more complicated input signal; by linearity, the channel output will simply be the superposition of the system response to each of the individual components. Since **S** contains information about N different frequencies, the output can be determined for a hypothetical input comprising components at all N of these frequencies.

One can use phasor notation to rewrite the channel input signal $x(t) = A\cos(2\pi f_L t+\theta)$ as $X = A\exp(j\theta)$. (The lower-case letter indicates the time-domain form, the capital indicates the phasor form.) Then the channel output corresponding to this input is found by a simple complex multiplication: $Y(\phi) = A\exp(j\theta) \times S21(\phi, f_L)$. A superposition of N sinusoids at N different frequencies can be written as follows, in time-domain form or as an N-element vector:

$$x(t) = \sum_{k=0}^{N-1} A_k \cdot \cos[2\pi(f_L + k\Delta f)t + \theta_k] \tag{3.16}$$

$$X = [A_0\exp(j\theta_0)\ \ A_1\exp(j\theta_1)\ \ \ldots\ \ A_{N-1}\exp(j\theta_{N-1})]^T. \tag{3.17}$$

Obviously, the first element in X corresponds to frequency f_L, the second to $f_L+\Delta f$, and so on. The channel output Y that corresponds to this input signal is expressed in this vector-phasor notation simply by multiplying each element of X by the appropriate S parameter:

$$Y(\phi) = [S21(\phi, f_L)X_1\ \ S21(\phi, f_L+\Delta f)X_2\ \ \ldots\ \ S21(\phi, f_U)X_N]^T, \tag{3.18}$$

where X_n indicates the nth element of vector X, $1 \le n \le N$. This can be written in matrix form by defining a simple N-dimensional linear transformation, $\mathbf{T}(\phi)$, that transforms channel inputs to channel outputs:

$$\mathbf{T}(\phi) = \begin{pmatrix} S21(\phi, f_L) & 0 & \cdots & 0 \\ 0 & S21(\phi, f_L+\Delta f) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & S21(\phi, f_U) \end{pmatrix} = \begin{pmatrix} s_{q,1} & 0 & \ldots & 0 \\ 0 & s_{q,2} & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & s_{q,N} \end{pmatrix}, \tag{3.19}$$

where $q = 1+\phi/\Delta\phi$. Thus the channel output Y corresponding to input X is

$$Y(\phi) = \mathbf{T}(\phi)\, X.$$

The time-domain form of $Y(\phi)$ is

$$y(\phi, t) = \mathrm{Re}\left\{ \sum_{k=0}^{N-1} Y_{k+1}(\phi) \cdot \exp\left[j2\pi(f_L + k\Delta f)t\right] \right\} . \qquad (3.20)$$

The same formula can be used to compute x(t) given X. This is simply an adaptation of the well-known formula for transforming phasor quantities into the time domain, found in [Stu81].

Let D denote the domain of this transformation. Clearly from the above discussion, a signal in D must have a power spectrum that does not occupy any frequencies other than the N sampled frequencies between f_L and f_U. Input signals with power outside of these discrete frequencies are outside the domain of $\mathbf{T}(\phi)$; the corresponding outputs cannot be computed. That is, the transformation will only predict the outputs to those inputs that fit in the spectrum of the original measurement. This was illustrated in Fig. 1-6, where the spectrum of the synthetic pulse excitation is aligned, in the frequency domain, with that of the channel measurements.

In the discussion of Sec. 3.2, the IFT of the S21 measurement was shown to be the channel response to an input r(t) that approximates an impulse; the corresponding phasor input R is uniform-amplitude excitation across the measurement bandwidth. When **S** is collected at N discrete frequencies,

$$R = [1\ 1\ \ldots\ 1]^T, \qquad (3.21)$$

$$r(t) = \mathrm{Re}\left\{ \sum_{k=0}^{N-1} 1 \cdot \exp\left[j2\pi(f_L + k\Delta f)t\right] \right\} \qquad (3.22)$$

$$= \sum_{k=0}^{N-1} \cos\left[2\pi(f_L + k\Delta f)t\right] \qquad (3.23)$$

$$= \cos(2\pi f_c t) \cdot \frac{\sin(\pi t \Delta f N)}{\sin(\pi t \Delta f)} , \qquad (3.24)$$

where $f_c = f_L + \frac{1}{2}\Delta f\,(N-1)$, which is midway between f_L and f_U. Note that r(t) is a periodic signal--a sinusoidal carrier modulated by periodic sinc-shaped pulses. The discrete-

frequency nature of R makes r(t) periodic. The derivation of (3.24) from (3.23) is found in Appendix B. Fig. 3-4 shows what this signal looks like in the time domain. Thus we see that the transmission coefficient measurement **S** reveals the channel output to a particular input signal, R, which is one of an infinite number of input signals for which the channel output can be determined; or, R is one member of D, the domain of the transformation $\mathbf{T}(\phi)$, and D is an infinite set (namely, N-dimensional complex space).

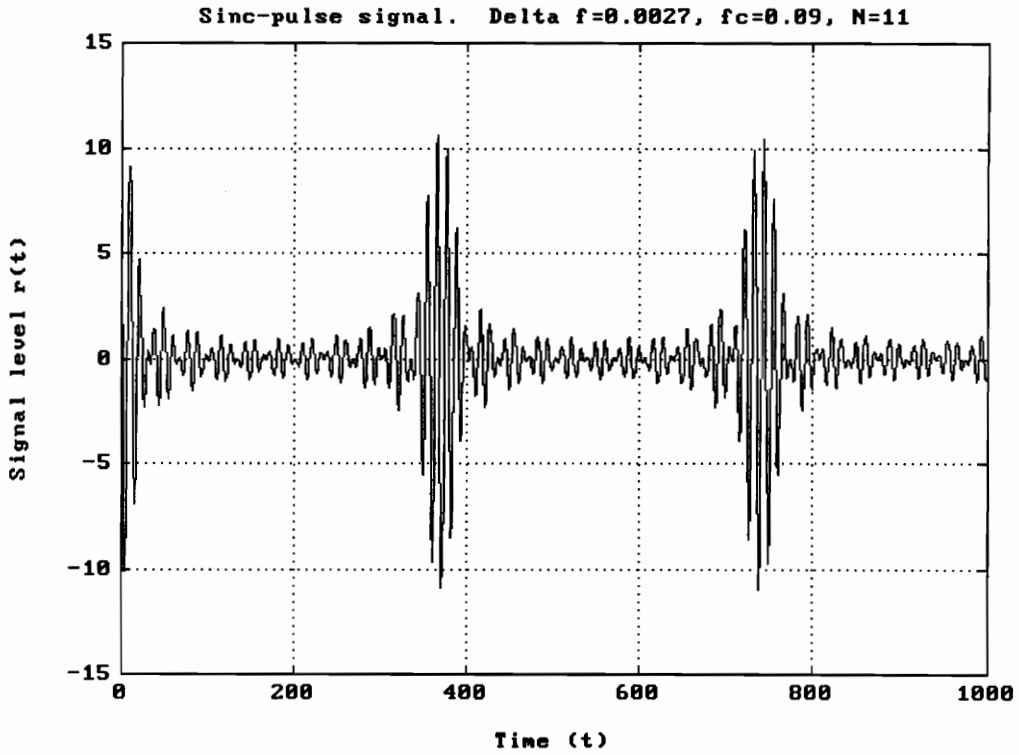

Figure 3-4: Example of r(t) when $\Delta f = 0.0027$, f_c=0.09, N=11.

Note that most of the members of D are purely hypothetical signals--superpositions of some or all of the N test frequencies--and cannot easily be physically realized. However, as long as the channel is linear, the channel output nevertheless can be determined for any member of D; that is, if one of the signals in D *were* synthesized (with the appropriate, expensive RF hardware), and used to excite the channel, the output would be almost exactly that predicted by the transformation $\mathbf{T}(\phi)$; the difference would

be caused by factors such as noise, intermodulation in the channel, and other real-world phenomena. For this reason these inputs and outputs are called synthetic. If pulse-shaped signals are considered, they are called synthetic pulses.

Of course, at this point, gating--that is, removal of interference--has not yet been considered; the focus has been how to determine the channel response to certain signals. The reason synthetic pulses merit attention is their application to gating. Ordinarily R of (3.21) is not the best choice for gating; other more promising synthetic pulses in D will be discussed below. Note also that there are perhaps members of D that are not pulse-shaped (therefore useless for gating) that are appropriate to other interference-removal algorithms. For example, in Chap. 6, application of spread-spectrum interference cancellation techniques will be considered very briefly; some signals in D could be appropriate for those interference cancellation techniques.

3.4.2 Use Of The IDFT

As we have seen, it is very simple to compute the output to a synthetic input with phasors via the matrix $\mathbf{T}(\phi)$. However, in order to visualize the presence of multipath components, it is necessary to show the signals (both input and output) in the time domain. This involves (3.20). A typical synthetic input is an RF pulse with a relatively high-frequency carrier component times a narrowband envelope, like DSB-SC modulation [Cou90], as seen in Fig. 3-4. Note that it is the envelope of the output signal that reveals the presence of multipath, whereas the carrier frequency does not change. Since the carrier component, apparent in Fig. 3-4, is distracting, it is advantageous to "demodulate" the output, i.e., remove the carrier component. Since $y(\phi, t)$ is a real-valued bandpass signal, centered in the frequency domain at f_c, we can write it in the form

$$y(\phi, t) = \mathrm{Re}\{ \exp(j2\pi f_c t)\, y_{CBB}(\phi, t) \} \tag{3.25}$$

$$= \cos(2\pi f_c t)\, y_{IBB}(\phi, t) - \sin(2\pi f_c t)\, y_{QBB}(\phi, t) \tag{3.26}$$

where y_{CBB} is a complex-valued baseband signal, with real and imaginary parts y_{IBB} and y_{QBB}. The exponential term represents the carrier frequency. We remove it by downconverting and filtering y, as follows, with LPF{} indicating ideal lowpass filtering:

$$y_{CBB}(\phi, t) = \text{LPF}\ \{\ \exp(-j2\pi f_c t)\ y(\phi, t)\ \} \tag{3.27}$$

$$= \exp(-j2\pi f_c t) \sum_{k=0}^{N-1} Y_{k+1}(\phi) \cdot \exp[j2\pi(f_L + k\Delta f)t]. \tag{3.28}$$

Eq. (3.28) follows from (3.27) because the spectrum defined by the phasors in Y is single-sided and multiplication by a complex exponential at frequency $-f_c$ shifts the spectrum directly to baseband: no double-frequency terms are generated, and the filtering is not needed. Simplifying,

$$y_{CBB}(\phi, t) = \sum_{k=0}^{N-1} Y_{k+1}(\phi) \cdot \exp[j2\pi(f_L + k\Delta f - f_c)t] \tag{3.29}$$

$$= \sum_{k=0}^{N-1} Y_{k+1}(\phi) \cdot \exp\left[j2\pi\Delta f\left(k - \frac{N-1}{2}\right)t\right] \tag{3.30}$$

Next we change the index from k to m where m = k - (N-1)/2; then we compute samples of y_{CBB} at time 0, T, 2T, and so on, where $T = (\Delta f P)^{-1}$, and P, the pad value, is a positive integer. This is one reason why it is convenient that N be odd: we want m to be an integer.

$$y_{CBB}(\phi, t) = \sum_{m=-(N-1)/2}^{(N-1)/2} Y_{m+(N+1)/2}(\phi) \cdot \exp(j2\pi\Delta f m t) \tag{3.31}$$

$$y_{CBB}(\phi, nT) = \sum_{m=-(N-1)/2}^{(N-1)/2} Y_{m+(N+1)/2}(\phi) \cdot \exp(j2\pi mn/P) \tag{3.32}$$

Except for a scaling factor of 1/P, this can be recognized as the well-known inverse discrete Fourier series (IDFS) of Y [Opp89]. The DFS and IDFS relate discrete-time periodic signals to their frequency domain (or phasor domain) representation. This can

worked into a more convenient form by breaking the summation into positive-index and negative-index parts:

$$y_{CBB}(\phi, nT) = \sum_{m=-(N-1)/2}^{-1} Y_{m+(N+1)/2}(\phi) \cdot \exp(j2\pi mn/P) + \sum_{m=0}^{(N-1)/2} Y_{m+(N+1)/2}(\phi) \cdot \exp(j2\pi mn/P) \tag{3.33}$$

$$= \sum_{m=-(N-1)/2}^{-1} Y_{m+(N+1)/2}(\phi) \cdot \exp[j2\pi(m+P)n/P] + \sum_{m=0}^{(N-1)/2} Y_{m+(N+1)/2}(\phi) \cdot \exp(j2\pi mn/P). \tag{3.34}$$

Since n is an integer, in (3.34), $\exp(j2\pi mn/P) = \exp[j2\pi(m+P)n/P]$. We now change variables again, replacing m in the left summation by p, where p = m+P. Therefore,

$$y_{CBB}(\phi, nT) = \sum_{p=P-(N-1)/2}^{P-1} Y_{p-P+(N+1)/2}(\phi) \cdot \exp(j2\pi pn/P) + \sum_{m=0}^{(N-1)/2} Y_{m+(N+1)/2}(\phi) \cdot \exp(j2\pi mn/P). \tag{3.35}$$

Note the similar exponential terms in the two terms above. Except for a scaling factor of 1/P, these two summations represent two parts of an inverse discrete Fourier transform (IDFT) on frequency-domain data zero-padded to length P, assuming P is larger than N. If $\log_2 P$ is an integer, then P samples of y_{CBB} can be computed efficiently using an inverse fast Fourier transform (IFFT) algorithm. The magnitude of y_{CBB} then reveals the amplitude of the excitation incident on the AUT.

If we use an IFFT of length P to compute the samples of y_{CBB}, then the input to the IFFT is an array or vector of length P. The first (N+1)/2 elements of the input vector are the last (N+1)/2 elements of Y, and the last (N-1)/2 elements of the input vector are the first (N-1)/2 elements of Y. The middle P-N elements are zero. As an example, consider this computation at a fixed azimuth for which P=16 and N=5. The input vector would be

$$W = [Y_3\ Y_4\ Y_5\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ Y_1\ Y_2]. \tag{3.37}$$

Then if one computes w=IFFT(W), using any of several popular math packages such as MATLAB, w will be filled with sixteen samples of the complex envelope corresponding to synthetic output Y. The samples will be spaced $T=0.0625/\Delta f$ apart in time, and scaled by 0.0625:

$$w = [y_{CBB}(0)\ \ y_{CBB}(T)\ \ y_{CBB}(2T) \ldots y_{CBB}(15T)] \cdot 0.0625. \qquad (3.38)$$

Sampling beyond time 15T is superfluous, because y_{CBB} is periodic with period $1/\Delta f$ (see Appendix B). If finer time-sampling is desired, P should be increased. The scaling factor of 1/P is not of concern to the user interested in inspecting y_{CBB}, since in practice y_{CBB} is displayed on screen normalized to a peak of 0 dB; however, the scaling factor of 1/P is important for reasons to be discussed in Sec. 3.5.2.

It is just as simple to use the IDFT to compute time-domain samples of synthetic inputs; (3.35) still applies, and an IFFT may be used in a manner parallel to that of the above example; in this case the elements in the IFFT input array will be the channel input vector of phasors, as in (3.17), and the IFFT result will contain samples of the time-domain input, as in (3.16). In the next section we discuss various choices for the phasor input vector.

3.4.3 Frequency-Domain Tapering

As mentioned, the S21 channel measurements directly yield the channel response to a synthetic sinc-pulse input signal described by (3.24). However, other synthetic inputs are worth considering. A very useful class of synthetic inputs is defined by standard tapering functions used in many engineering applications. The four considered here, plotted in Fig. 3-5 (a), are [Opp89]

Rectangular: $w_R(n) = 1, \qquad 0 \le n < N \qquad (3.39)$

Hann: $w_N(n) = 0.50 - 0.50 \cos [2\pi n/(N-1)], \qquad 0 \le n < N \qquad (3.40)$

Hamming: $w_M(n) = 0.54 - 0.46 \cos [2\pi n/(N-1)], \quad 0 \le n < N$ (3.41)

Blackman: $w_B(n) = 0.42 - 0.50 \cos [2\pi n/(N-1)] + 0.08 \cos [4\pi n/(N-1)], \quad 0 \le n < N$ (3.42)

These functions are all zero for $n < 0$ and $n \ge N$. They are used as the amplitudes of the phasor elements in the input vector X defined in (3.17). The phase term of each phasor element is zero for these inputs. Each defines a particular pulse shape for the channel. For completeness, (3.39) defines the "rectangular taper" corresponding to sinc-pulse input R of (3.21) - (3.24), even though it is arguably not a taper at all. The other three equations define pulse shapes that differ from r(t). For example, one might choose to use a Hamming taper to define the channel input. Then the phasor-domain synthetic input X and output Y are

$$X = [w_M(0) \;\; w_M(1) \ldots w_M(N-1)]^T \quad (3.43)$$

$$Y(\phi) = \mathbf{T}(\phi)\, X \quad (3.44)$$

$$= [S21(\phi, f_L)\, w_M(0) \;\; S21(\phi, f_L+\Delta f)\, w_M(1) \;\; \ldots S21(\phi, f_U)\, w_M(N-1)]^T. \quad (3.45)$$

Any of the windows in (3.39) to (3.42) can be used this way. Then one can use the IFFT method presented above to generate time-domain samples of X and Y.

The advantage of using one of the tapers in (3.40) - (3.42) as a synthetic input is that the high time-domain sidelobe level (ringing) of r(t) is reduced. However, the width of the main pulse increases as ringing decreases. Other properties, such as the pulse repetition rate, remain the same, since they are determined by parameters such as N, B, and Δf, not by the choice of synthetic input. Fig. 3-5 (a) illustrates these frequency-domain phasor weighting for the above tapers, and Fig. 3-5 (b) shows the corresponding time-domain pulse envelope, using a normalized decibel scale. The time axis of these plots are normalized to one pulse period, that is, $1/\Delta f$. One can clearly see from the upper left corner of Fig. 3-5 (b) that a rectangular-tapered (that is, untapered) phasor input yields a pulse envelope with the narrowest main pulse of the above four weightings. This, of

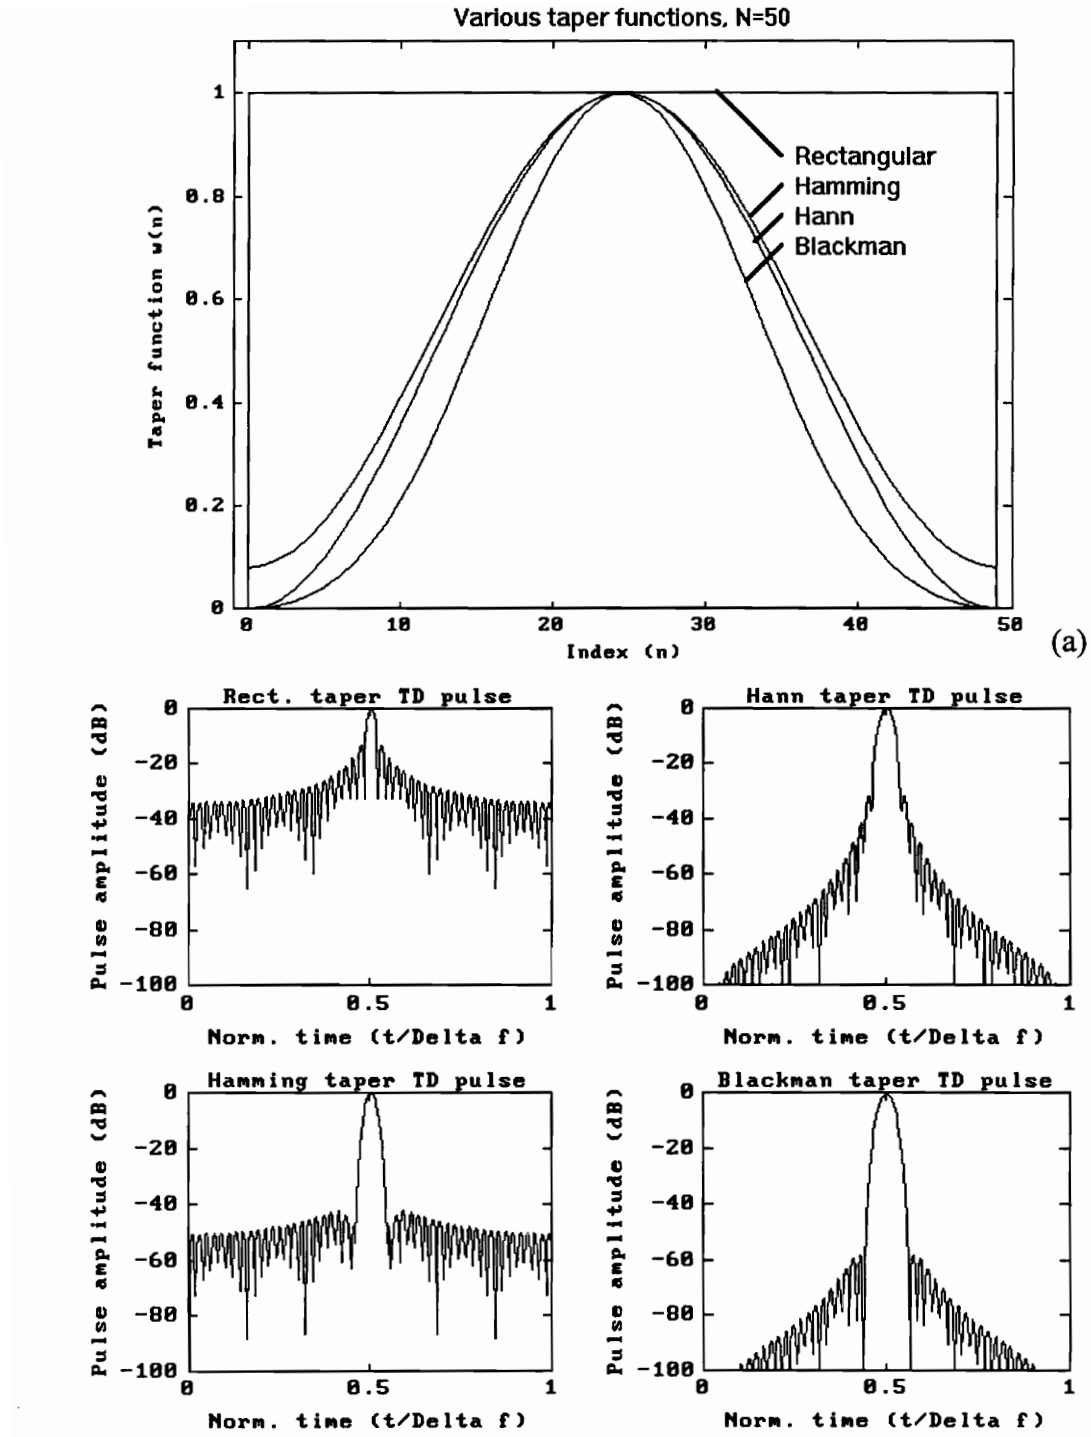

Figure 3-5: Illustration of frequency-domain tapers. (a) FD tapers (phasor weights). (b) TD synthetic pulse envelope corresponding to each FD taper.

course, is the weighting of R, shown in (3.21). The Hamming taper yields the next-narrowest pulse, the Hann taper yields a still-wider pulse, and the Blackman taper yields the widest pulse. The ringing levels parallel this trend: the rectangular taper causes peak ringing only 13 dB below the main pulse, and little decrease in ringing with increasing time. The other pulse shapes cause lower ringing and comparable (Hamming) or decreasing (Hann, Blackman) ringing versus time. The ringing level is important for visually identifying sources of reflections, since reflection levels may be many decibels below the amplitude of the line-of-sight component, and cannot be observed if they are masked by the ringing of a nearby component.

3.5 Gating Theory

Once the line-of-sight and interference components have been visually identified from the time-domain plot of the synthetic output $y(\phi, t)$, the next step is to apply an appropriate time gate $g(t)$ to preserve the line-of-sight component and attenuate the interference. Fig. 4-8 (b) and (d) show an example time domain plot and a time gate. Once the gate is chosen, the gated time-domain data $z(\phi, t)$ are transformed back to the frequency domain in order to obtain frequency-domain transmission coefficient data $Z(\phi, f)$ that, ideally, pertain solely to the line-of-sight path. This new transmission coefficient data should yield pattern measurements via (3.1) more representative of the AUT's true pattern than do the ungated S21 data $Y(\phi, f)$.

Gate functions were illustrated in Fig. 1-7. The gate width, position, and shape are all parameters under the user's control. As will be discussed, a well-designed gate will be neither too narrow nor too wide. In certain cases (e.g., testing a very narrowband antenna) the interference and line-of-sight components will be indistinguishable, since the AUT will smear the line-of-sight component and interference components together, and no gate function can remove the contamination. Unfortunately for the user, and despite the efforts discussed in Sec. 2.5, there is no systematic way to choose an appropriate gate

function, and the user of the system will need judgment and discretion to remove interference via gating. The guidelines for the gate design are presented later.

The basic procedure is simple. The user needs to center the gate on the line-of-sight component, to make the gate wide enough to pass the line-of-sight component, and to make the gate narrow enough to exclude the multipath effects. Note that gate design will generally involve compromise: some short-delay range multipath might need to be tolerated in order to prevent pattern distortion due to inappropriate gating. That is, it would be better to choose a gate that suppresses multipath by, say, 10 dB, and preserves the integrity of the pattern, instead of a narrow gate that suppresses multipath by 30 dB while simultaneously causing significant pattern errors. As an example of the distortion, Evans [Eva90a] shows that while measuring sidelobe levels in a reflector antenna, a too-narrow gate will cause the sidelobes to appear artificially low. This subject will be discussed further in Sec. 3.5.3 and 3.5.4.

3.5.1 Basic Operation And Application Of The DFT

Once a gate g(t) has been designed, it is multiplied with the complex synthetic output, y_{CBB}, from (3.35), to form a gated result z_{CBB} in the time domain. Then this complex envelope is used to modulate a carrier just as in (3.25):

$$z_{CBB}(\phi, nT) = g(nT) \cdot y_{CBB}(\phi, nT) \tag{3.46}$$

$$z(\phi, t) = \mathrm{Re}\{ \exp(j2\pi f_c t)\, z_{CBB}(\phi, t) \}. \tag{3.47}$$

The gate has a peak value of unity so that it preserves the gain of the line-of-sight transmission coefficient. Note that g(t), like y_{CBB}, is periodic in the time domain with period $1/\Delta f$; thus z_{CBB} is periodic with the same period, and can be expressed by a Fourier series; in other words, it is a discrete-frequency signal with a fundamental frequency of Δf. Strictly speaking, z_{CBB} has infinite bandwidth since g will generally have abrupt stop-intervals; nevertheless we will only consider a bandwidth of B; if the gate is chosen

appropriately, the out-of-band power will be small. If the gate is made too small, the out-of-band power could be significant. We will return to this subject in Sec. 3.5.3.

Since the channel output y(t) is multiplied in the time domain with a gate function g(t), the gating system can be thought of as a time-varying system. As a result, z(t) can comprise components at frequencies other than the N frequencies between f_L and f_U. Lacking a better name, we will refer to these components as modulation products. Compared to out-of-band power due to gating, modulation products will normally be very small. Thus we will henceforth ignore their contribution, and restrict our attention to the components of z(t) at f_L, $f_L+\Delta f$, . . . , f_U, which we will assemble into an N-element vector Z of phasors just like the vectors X and Y of the previous sections. To determine the elements of Z from the time-domain samples of z_{CBB} we need a procedure the reverse of the one in Sec. 3.4.2 that determined samples of y_{CBB} from Y. As shown in [Opp89], the discrete Fourier transform (DFT) is the complement of the IDFT, and will compute the desired DFS coefficients (phasors) that are the elements of Z. This can be demonstrated by substituting Z_k and z_{CBB} for Y_k and y_{CBB} into (3.32) and then solving for Z_k, in the manner shown in [Opp89]:

$$Z_k(\phi) = \frac{1}{P}\sum_{n=0}^{P-1} z_{CBB}(\phi, nT)\exp\left[-j2\pi\left(k-\frac{N+1}{2}\right)n/P\right], \quad 1 \le k \le N. \quad (3.48)$$

For values of k between (N+1)/2 and N, this is simply the DFT of z_{CBB} scaled by 1/P,

$$Z_k(\phi) = \left(\frac{1}{P}\right) V_{k-(N+1)/2}, \quad \frac{N+1}{2} \le k \le N, \quad (3.49)$$

where V_r is the rth element of the DFT of z_{CBB}. When k is between 1 and (N-1)/2, this is equivalent to the following, due to the periodic nature of the DFT:

$$Z_k(\phi) = \left(\frac{1}{P}\right) V_{k+P-(N+1)/2}(\phi) = \frac{1}{P}\sum_{n=0}^{P-1} z_{CBB}(\phi, nT)\, \exp\left[-j2\pi\left(k+P-\frac{N+1}{2}\right)n/P\right],$$
$$1 \le k \le (N-1)/2. \quad (3.50)$$

Note the scaling factor of 1/P. This factor is conventionally part of the IDFT but absent from (3.35), and it is conventionally absent from the DFT but appears here. Of course, there is no advantage to multiply the IDFT output by P, then the DFT output by 1/P, since the transforms are homogeneous; instead we use the IDFT and DFT to compute y_{CBB}/P and Z, respectively, and in effect apply the scaling factor before, instead of after, the gating.

Parallel to the earlier discussion, if $\log_2 P$ is an integer, then (3.49) and (3.50) can be computed efficiently using an FFT. For example, if N=5 and P=16, then we assemble a time-domain vector of z_{CBB} samples as follows:

$$w = [z_{CBB}(0)\ \ z_{CBB}(T)\ \ z_{CBB}(2T) \ldots z_{CBB}(15T)] \cdot 0.0625. \quad (3.51)$$

Then, using MATLAB, we compute W=FFT(w), which yields the following phasors:

$$W = [Z_3\ Z_4\ Z_5\ * * * * * * * * * * *\ Z_1\ Z_2]. \quad (3.52)$$

The asterisks denote out-of-band components; these values will generally be nonzero. The out-of-band terms close to Z_1 and Z_5 will tend to be larger than the terms in the middle of W. This introduces an important point that will be further developed below; for now, we need only mention that in particular circumstances these out-of-band components can become "folded" back into the phasors of Z, causing aliasing errors.

3.5.2 Time-Domain Tapering

Just as it is advantageous to taper the phasor input in order to suppress time-domain sidelobes (ringing), it is advantageous to taper the gate to suppress frequency-

domain sidelobes, which are directly related to the above-mentioned out-of-band components. The tapers under consideration are the same as those presented in Sec. 3.4.3, scaled to the appropriate gate length. Fig. 3-5 illustrated the relationship between the frequency-domain (phasor) taper and time-domain pulse envelope; a similar relationship holds between the time-domain (gate) taper and the frequency-domain smoothing function, which will be discussed in the next section.

3.5.3 Time-Domain Gating As Frequency-Domain Smoothing

So far, the gating procedure has been presented from an intuitive perspective; the gate should permit line-of-sight energy and inhibit interference. Unfortunately, this point of view gives little insight into all the consequences of the gating operation. By presenting gating from a different perspective, that is, from a frequency-domain perspective, we will establish three important considerations for gate design beyond the obvious ones already discussed.

Convolution in the time domain is equivalent to multiplication in the frequency domain and vice versa. Eq. (3.46) shows that the gate is applied to y_{CBB} in the time domain via multiplication, and therefore, in the frequency domain, Z represents the convolution of Y and another function, G. In the frequency domain, then, the gating function smoothes out Y just as, for example, the video filter on a spectrum analyzer smoothes out a noisy spectrum. This is illustrated in Fig. 1-7; in the upper left corner the spectrum of Y is rippled due to the superimposed interference component, but after time-domain gating, the spectrum of Z in the lower right corner is smooth, because in the frequency domain the gating acts just like convolution with a filter, G, which we shall give the descriptive name "smoothing function."

The first phenomenon we shall consider is the *edge effect*. That is, the phasor data at the edges of B (that is, near f_L and near f_U) will be attenuated: the corners, in the frequency domain, will be smoothed off. This is illustrated in Fig. 3-6. Observe that the

convolution of two rectangular shapes yields a trapezoidal shape for which the top is as wide as the difference in width between the two components of the convolution. The same effect happens when the pattern data are gated; B represents the bandwidth of the phasor data, and W_G represents the bandwidth of the smoothing function. The gated results are artificially attenuated at f_L and f_U. Therefore the reference levels (cf. Sec. 3.1) of the patterns near f_L and f_U will not be meaningful; the AUT will appear to have lower gain near these frequencies. The only meaningful pattern reference levels will occur for the patterns in the middle of the measurement bandwidth, in the part of the band analogous to C in Fig. 3-6. (This effect is not always a matter of concern, since the user will often be interested only in the shape of the pattern, and uninterested in the reference level that relates different patterns.) So, if the user needs patterns across a bandwidth C for which the reference levels are comparable to each other, then B must exceed C. The amount of excess bandwidth will depend upon the width and taper of the time-domain gate, according to this approximate formula:

$$C = B - 4/T_G. \tag{3.53}$$

The units of C and B are hertz and the gate width, T_G, is in seconds. This is an approximate formula based directly on Figs. 3-5 and 3-6, so we will not quibble about how best to define the width of the gate. As an example, suppose the user needs pattern data across 250 MHz centered around f_c for which the gain of the patterns can be compared to each other. Thus C = 250 MHz. One strategy is to make B = 500 MHz (centered at f_c, of course). Based on Fig. 3-5 (b), the synthetic pulses available will be about 8 ns wide, but according to (3.53) the gate function g(t) must be no narrower than $T_G = 4/(B-C) = 16$ ns in order to preserve the reference levels in the center 250 MHz of B. (Note that a 16 ns gate is not necessarily narrow enough to eliminate multipath--that is a separate consideration!) If 16 ns is too wide, another strategy is to increase B to 750 MHz; then the gate may be narrowed to about 8 ns, and once again the center 250 MHz will yield comparable

patterns. Unfortunately, the AUT usually dictates a maximum value of B--for example, an antenna with 5% bandwidth restricts B to a maximum of 0.05 f_c. If B cannot be increased, then the user gate must increase T_G to widen C. Again, this effect is not always a concern. If the user does not need to compare patterns of different frequencies, the gate may be as narrow as the AUT itself permits--which will be discussed next.

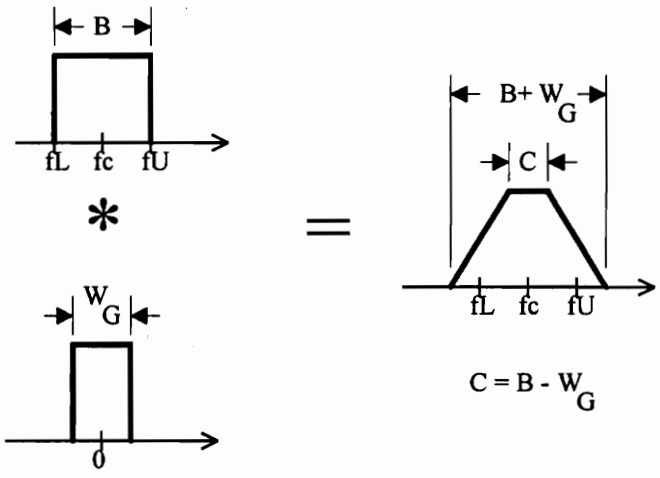

Figure 3-6: Convolution of rectangular-shaped signals.

The second frequency-domain effect of gating is the *flattening effect*. This effect has received most of the attention in the literature (cf. Sec. 2.5) because it is the hardest to evaluate. A narrow gate will smooth out the frequency-domain characteristics of the AUT response, just like a narrowband filter will smooth out the time-domain variations in a signal. This is illustrated in Fig. 4-8 (a) and (c), where the solid curves indicate the frequency response Z(f) of the gated channel response. The flattening effect can be useful, because range multipath manifests itself as a fast frequency-domain ripple in the data. Unfortunately a complication arises. One does not expect an antenna to have the same pattern at all frequencies; hence the frequency domain variations in the data are potentially an important part of the pattern measurement. If the gate is made too narrow, the smoothing function will smooth the frequency-domain variations of the true pattern,

which will cause pattern distortion, such as the artificial decrease in sidelobe levels mentioned in Sec. 3.5. Therefore it will generally be necessary to examine the AUT's response to the synthetic pulse and attempt to determine, via a visual inspection, the width of the line-of-sight response. Fig. 4-8 (b) and (d) show an example of how this is done. Then, to be safe, the user should make the gate somewhat wider than the desired response. Note that an AUT might be more narrowband in its sidelobes than in its main lobe; the time-domain response of a sidelobe would be longer than the response of the main lobe. Since the same gate is used for all azimuths, a user therefore needs to make the gate wide enough for the widest time-domain response of the AUT.

Another way to look at the flattening effect is to think of the gated pattern as a composite of patterns across a range of frequencies, rather than a CW pattern. Convolution, as is well known, can be though of as a weighted, moving average. The smoothing function G therefore causes the gated pattern at frequency f to represent a weighted average of the ungated patterns at frequencies neighboring f. The size of the "neighborhood" depends on the width of the gate, T_G. As stated above, the width of G is approximately $4/T_G$. When the gate is very wide (or not used at all), then the gated patterns represent the AUT response to very narrowband excitation, although interference will distort the shape. When the gate is narrowed, the gated pattern shows the AUT response to a wide-bandwidth stimulus, and is therefore not directly comparable to the pattern at any particular frequency. For example, suppose a 40 ns gate is applied to the channel output. This causes a gated pattern "at" frequency f actually to represent the average AUT response across a 100 MHz bandwidth centered on f. If the AUT is known a priori to be wideband, then this average perhaps reveals the CW pattern at f much better than the raw pattern measurements do. If, on the other hand, the AUT is narrowband, then perhaps the gated pattern reveals nothing about the AUT's true CW pattern. Therefore in order to prevent pattern distortion, the user should choose a gate wide enough that the true antenna pattern is not expected to change significantly within the "composite bandwidth."

The third effect we shall discuss is *aliasing*. This is an error that is easily avoided. As we know, time-domain multiplication implies frequency-domain convolution. Also, convolution can be linear or circular, and circular convolution (i.e., periodic convolution) can lead to aliasing [Opp89]. Since we are using a DFT to compute Z, circular convolution is occurring, and its effects must be considered. The aliasing problem arises when the chosen pad value, P, is too small relative to N and the gate width, and when the gate is untapered. This is illustrated in Fig. 3-7. The smoothing function G, when convolved with the phasors in Y, will generate out-of-band terms. As long as these out-of-band terms stay outside the measurement bandwidth B (of eq. 3.9), we do not care about them. However, if G is too wide or P is too small, they can alias into the bandwidth B. In fact, there will always be some aliasing, because G will typically have an infinite bandwidth due to the frequency-domain sidelobes from the gate shape; however, if the time gate is wide enough, and tapered, and P is large enough, the aliasing will be

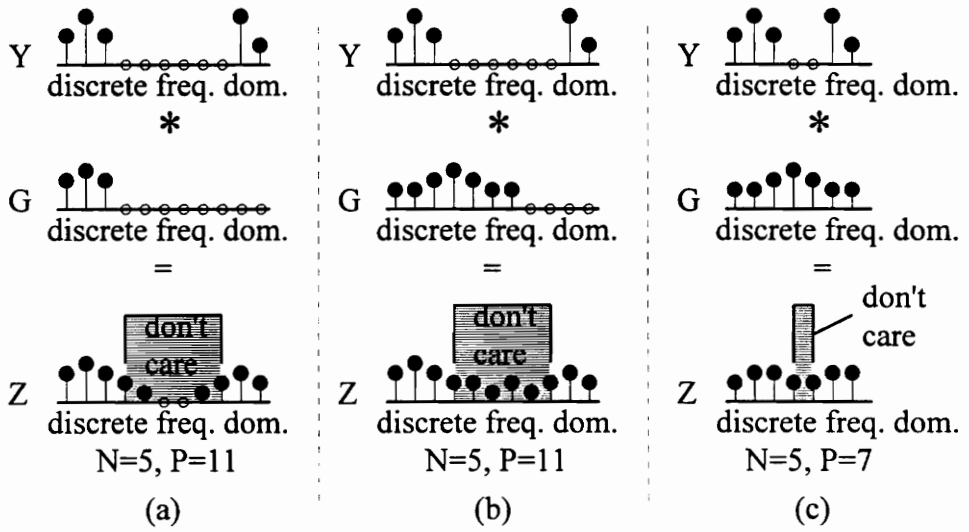

Figure 3-7: Illustration of three situations in frequency-domain circular convolution, Z = G * Y. (a) Ideal case: no aliasing in Z due to narrow G and large P. (b) Marginal case: no aliasing in the significant parts of Z due to large P. (c) Pathological case: aliasing due to wide G, small P.

negligible. The taper applied to the time gate causes G to be relatively restricted in bandwidth. If the gate is untapered (i.e., rectangular tapered) then G will have high frequency-domain sidelobes, which in effect widen G, even if the time gate itself is fairly wide. Based on these facts, then, as long as the time gate is at least as wide as the synthetic input pulse, then G will be no wider than B, and the bandwidth of Z will be no wider than 2B as seen from Fig. 3-6. Then a value of P at least twice N should keep the aliasing error negligible. So, the following is recommended:

$$P \geq 2N.$$

To summarize, the gate should be . . .

- . . . centered over the line-of-sight response in the time domain;
- . . . tapered, to limit the bandwidth of G;
- . . . wide enough to preserve the desired measurement bandwidth;
- . . . wide enough to pass the AUT's widest line-of-sight response;
- . . . wide enough to prevent aliasing; and
- . . . narrow enough to reduce interference.

Of the three phenomena considered in this section, the flattening effect is the one that requires the most judgment, since the true frequency-domain response of the AUT is almost never known a priori. The edge effect and aliasing effect are significant but they can be rendered benign by the appropriate selection of B, P and gate width, all of which are known quantities. Note that the last of these six objectives is in opposition to the prior three; the user will need to determine an acceptable compromise when deciding upon the best time gate for the application at hand.

3.5.4 Removal Of Frequency-Domain Taper

The frequency-domain smoothing function G ideally removes the effects of multipath. Once the gate is designed, this smoothing function does not depend on any particular synthetic input applied to the system; it can be applied directly to the set of transmission coefficient measurements, **S**, itself. This is desirable for reasons already discussed--the user might want to compare reference levels of patterns at different frequencies to see gain variations versus frequency for an antenna. If a frequency-domain taper remains attached to the pattern measurements, then the reference levels of the patterns are not directly comparable, and the gain variation is obscured. Therefore, once the user has decided on the time-domain gate that is appropriate to the data, the smoothing function is applied to the data in **S** without a frequency-domain taper (i.e., with a rectangular taper). The gated **S** array then yields pattern data at each of the N frequencies in B.

3.6 Conclusions

The principles of a software gating system were presented in this chapter. First we considered the issues involved in making transmission coefficient measurements on the channel central to the far-field range. Then we looked at the relationship between these measurements and antenna patterns, followed by a discussion of the effects of multipath on these transmission coefficient measurements. This led to the argument that the Fourier transform is useful in distinguishing the line-of-sight component from the interference components, under appropriate circumstances. Then we considered how a set of measurements **S** can be used as the basis for determining the channel output to an infinite set of synthetic inputs, how the synthetic inputs and outputs can be manipulated to reveal the presence of range multipath, and how that multipath can be removed by applying a time-domain gate. Finally, we considered how to design a time-domain gate, and how the gate affects the estimates of the line-of-sight transmission coefficients. In the next

chapter we will consider an application of all these principles: a software gating system implemented on the Virginia Tech antenna range.

CHAPTER 4

IMPLEMENTATION OF A SOFTWARE GATING SYSTEM

This chapter discusses a specific demonstration project that applies the principles of Chap. 3. This effort consisted largely of the development of software to collect and process antenna pattern data. The final result was an upgraded antenna range run by entirely new, easy-to-use software, with source code comprising over 100 original modules. The next chapter discusses the validation of the system.

4.1 Virginia Tech Rooftop Antenna Range Upgrade Project

The Antenna Group at Virginia Tech has had an antenna range on the roof of Whittemore Hall for several years. The range (pictured in Fig. 1-3) is illustrated schematically in Fig. 4-1. A source of signal (a sweeper or synthesizer) provides signal input to the source antenna. The AUT is mounted on the test tower, which is attached to an azimuth positioner. This positioner is turned by a DC motor controller, and its azimuth angle is measured via synchro windings, the signals from which are digitized and processed by a synchro-to-digital (SDC) converter. The AUT is attached to a receiver, which measures the channel output as the antenna turns. A data collection computer controls the test tower and stores the measurements.

Originally the channel output measurements were scalar (amplitude-only); however, an SA 1780 vector (amplitude and phase) measurement receiver was donated to the group by Bellcore in 1994. In early 1995 the group began an upgrade project with the following goals:

- Integrate the vector receiver into the pattern measurement system;
- Upgrade the signal source;
- Write new software, with an improved user interface, to collect pattern data; and
- Implement a software gating algorithm to process the pattern data.

Pattern distortion due to multipath, as illustrated in Fig. 1-2, has been recognized as a problem on the Virginia Tech antenna range for a long time. In the early stages of this project, the Antenna Group explored the utility of a software gating system for the range with an HP 8409 network analyzer. The HP 8409 was used to make stepped-frequency channel measurements, with the receive antenna azimuth fixed. Inspection of the channel measurements in the time domain indicated that the interference components were separable from the line-of-sight component for reasonable excitation bandwidths, which of course implies that, when measuring the patterns of the right types of antennas, the pattern distortion can be reduced. However, because of practical difficulties, it was impossible to adapt this system to measure antenna patterns. That is, with the HP 8409 it was possible to apply software gating techniques to a channel measurement at a fixed azimuth, but not possible to repeat these measurements at a series of azimuth steps and merge the data together. Therefore a new system was designed to accomplish this data collection and processing task.

The sweeper from the old system was unsuitable for because it does not possess sufficient frequency stability. Fortunately, at the time of the above investigation the group owned an HP 8673B, which operates from 2.0 to 26.0 GHz. To cover the low end of the spectrum, an HP 8648C was purchased that operates from 100 kHz to 3.2 GHz. The large overlap in frequency ranges helps prevent the situation in which one needs to use both synthesizers in a single measurement, since (at present) the source switchover from one synthesizer to the other is not computer-controlled.

The strategy chosen for the new software was to use the LabVIEW programming language by National Instruments. LabVIEW was selected because it makes graphical user interfaces easy to build, it is known as a language for reliable test instrument control, and--because it is promoted as an easy-to-learn language. It is hoped that the end product will be successfully maintained by students and engineers not otherwise at ease with programming.

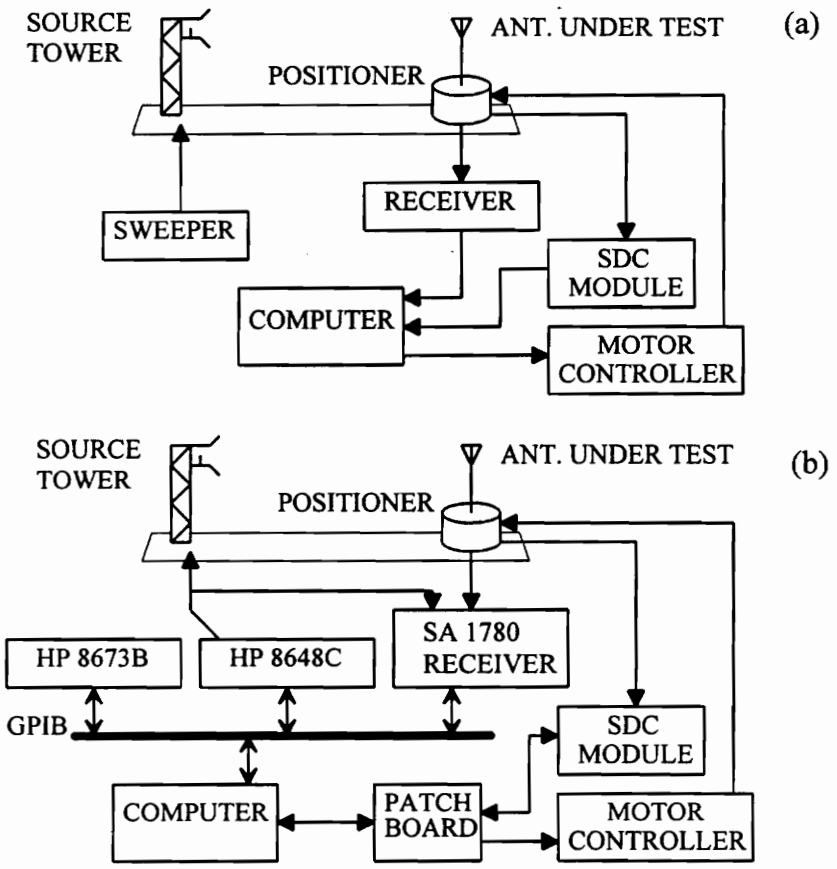

Figure 4-1: Antenna range architecture. (a) Old system. (b) Upgraded system.

4.2 Antenna Range Hardware

The hardware used in the old system is illustrated in Fig. 4-1 (a). The old data collection computer housed several digital input/output (DIO) cards which interfaced to

the receiver, the positioner's motor control, and the SDC via a custom digital interface. Unfortunately, these DIO cards were incompatible with LabVIEW, so a new LabVIEW-compatible DIO card was purchased. A patch board was built to connect the old SDC and motor controller lines to the new DIO card. The computer in the old system was an IBM-XT, and incapable of running LabVIEW, so a 80486 DX4/100 became the new system computer. The new receiver, as well as the new signal synthesizers, were GPIB compatible, so the computer was outfitted with a GPIB interface to control these instruments. The hardware of the final result is shown in Fig. 4-1 (b). The only major purchases for this project were the LabVIEW development system and the low frequency synthesizer; the other expensive items were already owned when the project began.

The DIO aspect of the range upgrade represents a small but significant part of the total effort. The upgrade can be built upon in the future. For example, the azimuth positioner contains two sets of synchro windings, on gearings of 1:1 and 36:1. Only the 1:1 windings are currently in use. Data from the 1:1 and 36:1 windings could be combined to give higher precision azimuth measurements, although this requires another SDC module to be integrated into the system. This sort of incremental upgrade would be nontrivial although not be very difficult, and presently there are many aspects of the range for which similar upgrades suggest themselves. Most of these upgrades involve interfacing with the data collection computer. Therefore it was of considerable concern to make the DIO interface design logical and expandable, to make it compatible with the old system yet ready for new features. Since compatibility and expandibility are somewhat at odds, and interface decisions of this type are almost irreversible, considerable time was spent by antenna group personnel in the DIO design decisions.

4.3 *Range Runner*

Development began on several fronts at once in February, 1995. The program design began with high-level planning of the program structure and sketches of the various

screens and user interface. The author learned LabVIEW using the Student Edition from Prentice-Hall, and built some drafts of the screens with this version while awaiting arrival of the full development system. The principles discussed in Chap. 3 were explored using MATLAB, and the review of the relevant literature began with [Eva90] and [Boy85]. Finally, interfacing issues were explored: the digital interface of the old system was investigated and documented more thoroughly, the DIO requirements for the new system were assessed, and the LabVIEW interface to the GPIB-controlled devices was begun.

The majority of the effort for this upgrade was spent by the author developing software. The resulting program, called *Range Runner*, performs several important tasks, including communication with the user, communication with the range instruments, data collection, data manipulation, and the software gating. From the start *Range Runner* was developed using a combination of top-down design, in the development of user interfaces and high-level program structure, as well as bottom-up design, in the development of the low-level modules such as the software instrument interfaces, the file formats, the program initialization routines, and certain aspects of the user interface (such as polar plots).

LabVIEW's reputed ease-of-use warrants brief comment. LabVIEW differs from programming languages like C and FORTRAN in two ways--the first superficial, the second profound. First, LabVIEW is a graphical programming language; that is, the source code produced by a programmer is essentially a diagram or schematic, rather than a listing of text. This is the basis for National Instruments' claim that it is easy to learn, although anyone who has worked with schematics knows that a complex diagram can be just as forbidding as a thick listing of text. Second, LabVIEW is a data-driven language [Nat94]; this means that the order of execution within the program depends upon the flow of data within the program. This is significantly different from traditional imperative languages [Seb89]; hence software development in LabVIEW requires a different mode of thought than does development in C or FORTRAN. From this experience the author concludes that although LabVIEW is not hard to learn, a programmer familiar only with traditional languages has little advantage over a bright non-programmer when learning

LabVIEW; and also, that writing good code in LabVIEW is essentially no easier or harder than writing good code in C or FORTRAN.

LabVIEW is very good for user and instrument interfacing. It certainly makes the development of a graphical user interface much easier than does a generic C compiler. LabVIEW provides a choice of graphical input devices--"buttons," "knobs," and so on, which the programmer can arrange on the screen, and (when the program runs) with which the user can easily specify data to the program. Likewise LabVIEW offers a variety of built-in graphs, indicators, and gauges with which the program can communicate data to the user. LabVIEW proved reliable in its communication with other instruments in the antenna range system (such as the receiver, motor controller, SDC, and synthesizers). Much of *Range Runner* dealt with prosaic affairs such as these. Fig. 4-2 shows an approximate breakdown of the applications of the 142 modules within version 1.0 of *Range Runner*. As this chart shows, the gating system itself represents only about 12% of the bulk of *Range Runner*. In terms of time, the first version of *Range Runner* required about eight man-months to complete.

An initial version of *Range Runner*, including the gating system, was complete in August 1995, when the author was assigned to a new project. However, throughout the fall the system was tested and occasionally refined, and the *Range Runner User's Guide* [Pre95] was completed in November. Chaps. 1 through 5 of this guide describe how to set up an experiment, collect the data, print the patterns, and save the data in a file. The gating system is described in Chap. 6 from a how-to perspective, and the system hardware and software are documented in Chaps. 7 and 8 for the benefit of those responsible for range maintenance.

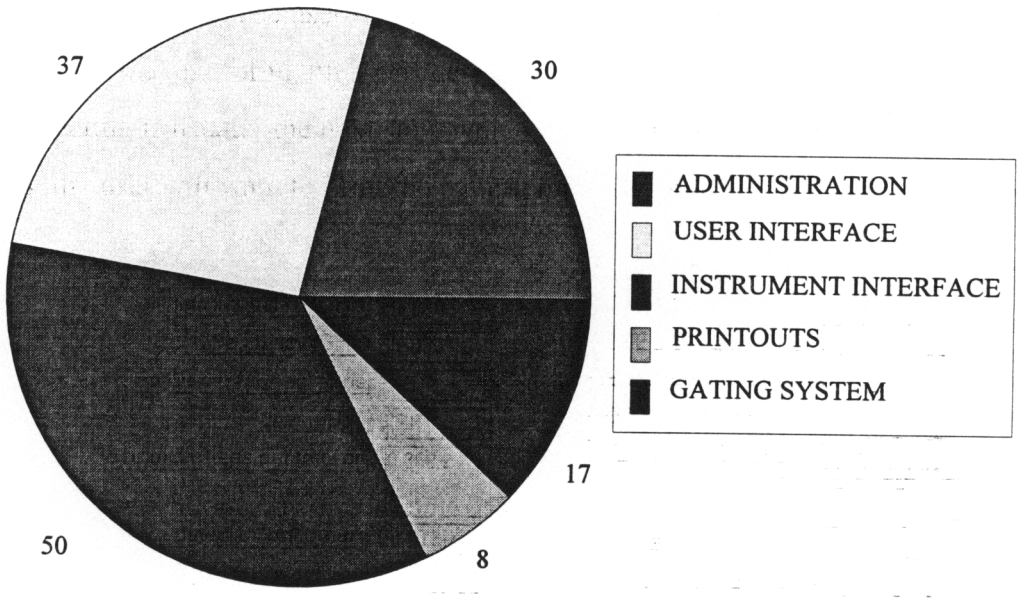

Figure 4-2: Categories of virtual instruments within *Range Runner*.

4.3.1 *Range Runner* Overview

In order to provide an overview of *Range Runner* we discuss the usual sequence of steps required for a pattern measurement on the antenna range. The following synopsis is an outline of how the program operates; although it omits many of the steps involved in making a measurement, it shows that the program addresses the pattern measurement requirements on the antenna range in a fairly comprehensive way. The first thing the user needs to do is to set up the source and test antennas, excite the source antenna with the appropriate synthesizer, and measure the channel output with the receiver. These instruments can be operated manually, and it is often very useful to manually verify that the receiver can detect a signal at the test antenna before starting *Range Runner*.

The first screen is main screen shown in Fig. 4-3. From this screen the user can load or save a data set, set up an experiment, perform an experiment, perform time-domain gating, print a pattern, or exit the program. The graph on the right shows the antenna pattern measurement (both ungated and gated patterns). Other controls on the

screen permit the user to adjust a few experiment parameters (the limits of the frequency sweep and the power level). At this point the user usually wants to set up an experiment, so he or she presses the CONFIG button, which brings up the configuration screen shown in Fig. 4-4.

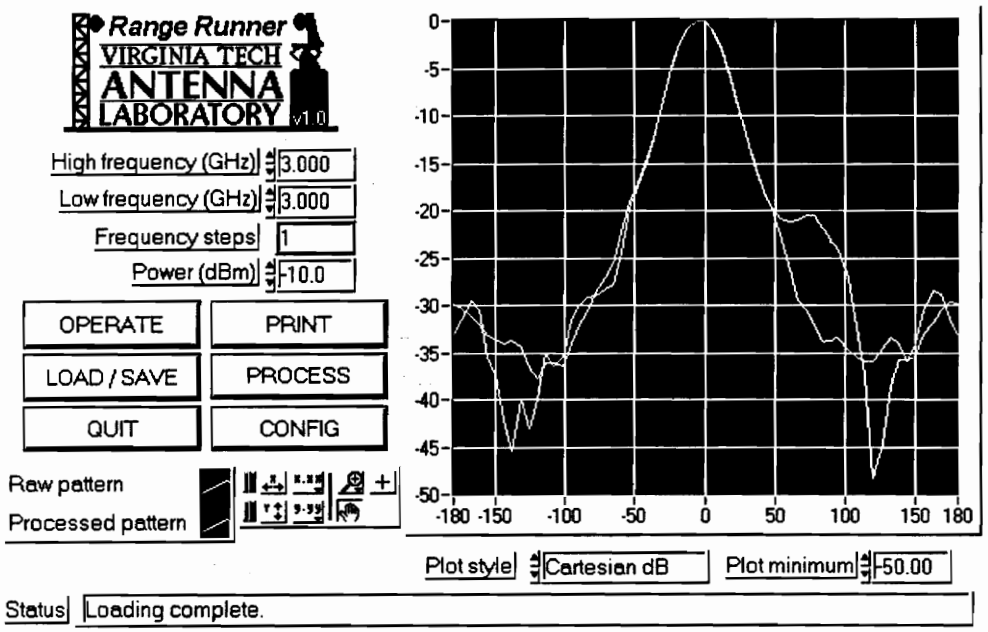

Figure 4-3: *Range Runner* main screen.

The configuration screen permits the user to control a large number of measurement parameters. The most basic are the frequency sweep specifications and the transmit power level. Other parameters include the IF bandwidth of the SA 1780 receiver and the postdetection averaging performed by the receiver. Useful information can be attached to the data set (such as transmitter polarization) even though this information does not affect the data collection. Default values for the gating can be specified from this screen, since *Range Runner* performs crude automatic gating using a few heuristics when the data collection is finished.

In the upper right corner of this screen (marked "For Your Information") *Range Runner* reports the approximate time required for the data collection--which can run to

several hours--and also the approximate pulse width and pulse repetition rate of the synthetic pulses, based on the user's chosen frequency sweep specification. It is good to know this information beforehand when one is relatively familiar with the characteristics of the antenna range, and when one knows about how narrow a pulse is necessary to distinguish range reflections.

The user can also specify the azimuth sampling in some detail via the control marked "Rx azimuth steps per rotation." The entry shown ("0, 360, 2") is reminiscent of a FORTRAN do-loop and here indicates that the pattern should be sampled in 2° azimuth increments from 0° to 360°. The start, end, and increment all can be specified by the user. Furthermore, the user can concatenate multiple sampling triplets with a semicolon, in order to sample certain parts of the pattern in greater detail. For example, if an antenna is expected to have nulls at 90° and -90°, one might want to sample the pattern more closely around the nulls. An appropriate azimuth specification might look like so:

0, 360, 2; 85, 95, 0.5; -85, -95, -0.5

This would sample the pattern in 2° azimuth increments except around the nulls, which would be sampled in 0.5° increments. A small lexical analyzer was written to implement this feature. Since each frequency sweep can be very time consuming, it is worthwhile to give this level of control to the user so that the minimum necessary number of pattern samples is measured.

Once the user has finished selecting experiment parameters, he or she can make the chosen parameters to be the *Range Runner* defaults for the next time the program is used, although normally this option is not chosen. This is effected by pressing ACCEPT. The user can select CANCEL to return to the main screen.

At this point the user is ready to press the OPERATE button on the main screen which initiates the pattern measurement process. Fig. 4-5 shows the screen that appears during the pattern measurement. The transmission coefficient measurements versus frequency are shown on the right hand side of the screen, and on the left a polar plot shows the ungated midband pattern as it is being measured. A large STOP button on this screen

permits the user to halt the measurement. Once the data collection is complete, the main screen will reappear. The user will generally want to save the data set at this point. The file is saved in a text format that was designed to be easily human-readable. Time-domain gating is usually the next step, which will be discussed below in more detail. After this, the user will want to print the pattern measurement.

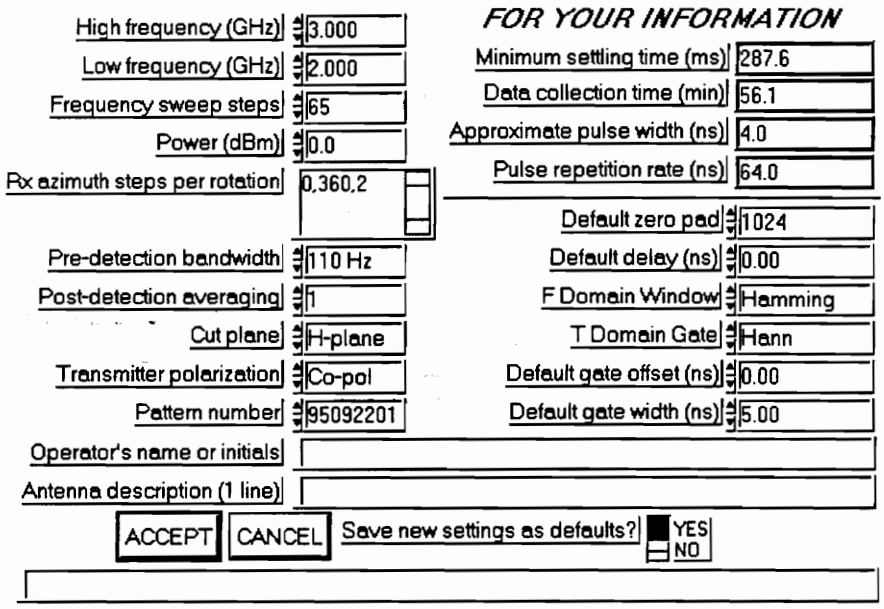

Figure 4-4: *Range Runner* configuration screen

Fig. 4-6 shows the printout setup screen, which offers the user a great deal of control over the appearance of the final printed result. The pattern itself can be shown in polar or Cartesian form, with a decibel or linear scale. The phase of the pattern can be printed as well. Each printout comes with a header at the top, which the user can manipulate. The user can select the frequency of the pattern to print, and choose between the ungated and the gated patterns.

Once the pattern has been printed and the data are saved to disk, the pattern measurement is complete, and the system can be turned off.

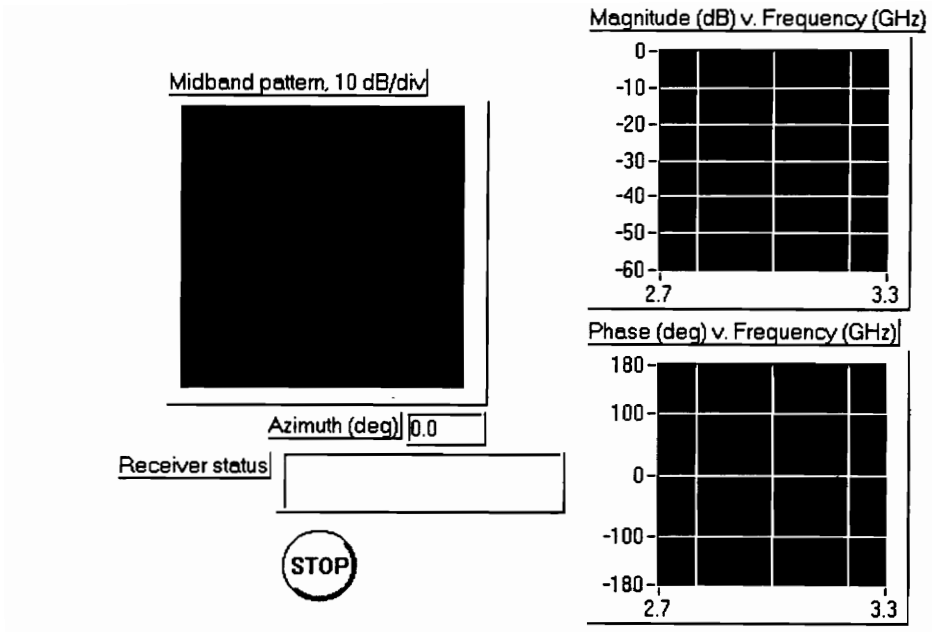

Figure 4-5: *Range Runner* pattern measurement screen

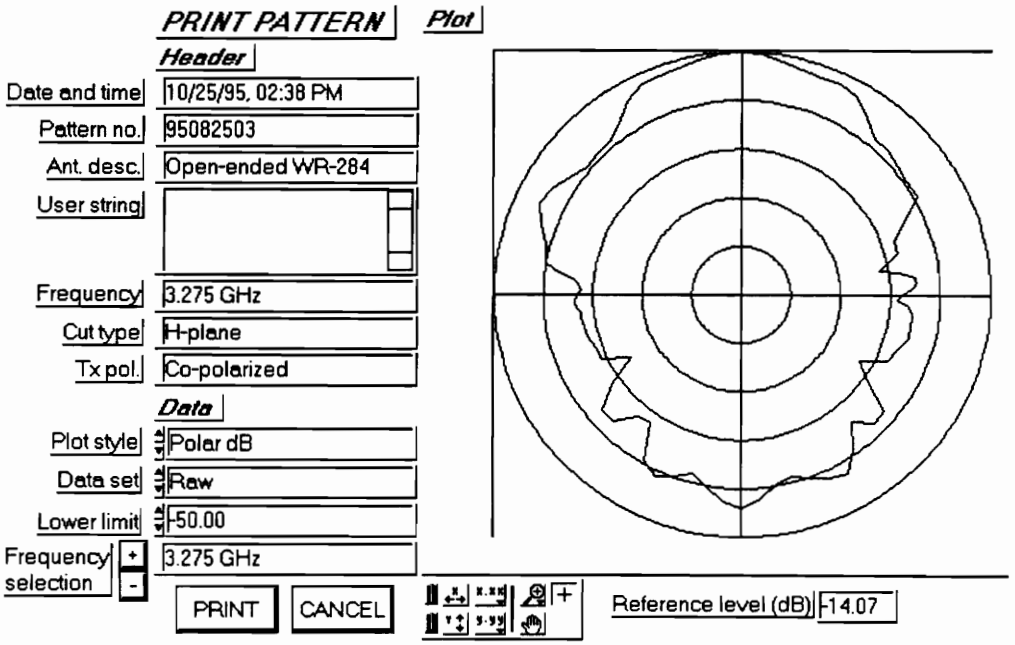

Figure 4-6: *Range Runner* printout setup screen

4.3.2 *GATELAB* Overview

The software gating portion of *Range Runner*, called *GATELAB,* is an interactive processing package which gives the user control over frequency-domain and time domain tapering, DFT pad size, gate location and width, and delay compensation. The *GATELAB* screen is shown in Fig. 4-7. This screen shows the stepped-frequency channel data for a particular AUT azimuth angle (which is chosen with a separate requester, and can be changed by pressing the NEW ANGLE button at the bottom of the screen). The channel data can be viewed in the frequency domain or time domain--the user chooses which via the control at the top of the screen. The time-domain signal is shown as an envelope on a decibel scale, so that low-amplitude interference is seen more easily.

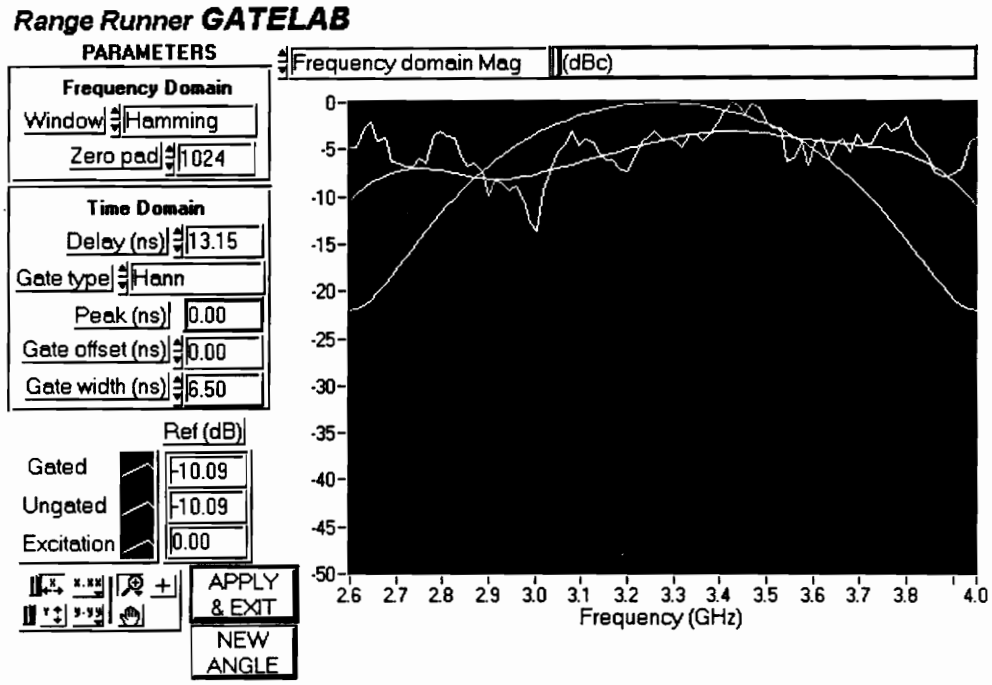

Figure 4-7: *GATELAB* screen

The curves on this screen are color-coded, although of course that is not apparent in the black-and-white plots presented here. The curve labeled "Excitation" shows the synthetic input to the channel. In frequency-domain mode, then, this is simply the

frequency-domain taper that the user has chosen, X(f); it is the dome shape in Fig. 4-7, which here is a Hamming taper. When the display is switched to time-domain mode, the excitation curve is the synthetic input envelope of x(t), showing the hypothetical channel input that would yield the "ungated" channel output shown. In frequency-domain mode, the curve labeled "Ungated" shows the raw S21 measurements T(f) at the selected azimuth. This corresponds to the jagged line in Fig. 4-7. In the time domain, the Ungated curve shows the channel's synthetic output y(t). The curve labeled "Gated" shows the data that are the result of the gating procedure. In time-domain mode, this is z(t), the product of the channel output and the gate. In frequency-domain mode, this is Z(f), the smoothed transmission coefficient spectrum; it appears in Fig. 4-7 as the smooth curve that roughly follows the peaks and dips of the ungated spectrum. For the reasons discussed in Sec. 3.5.4, the Gated curve in frequency-domain mode is smoothed because it is gated, but it is not tapered with a frequency-domain taper. This is apparent from Fig. 4-7, since clearly the Gated frequency domain curve does not show the 22 dB attenuation at the band edges that is characteristic of a Hamming window.

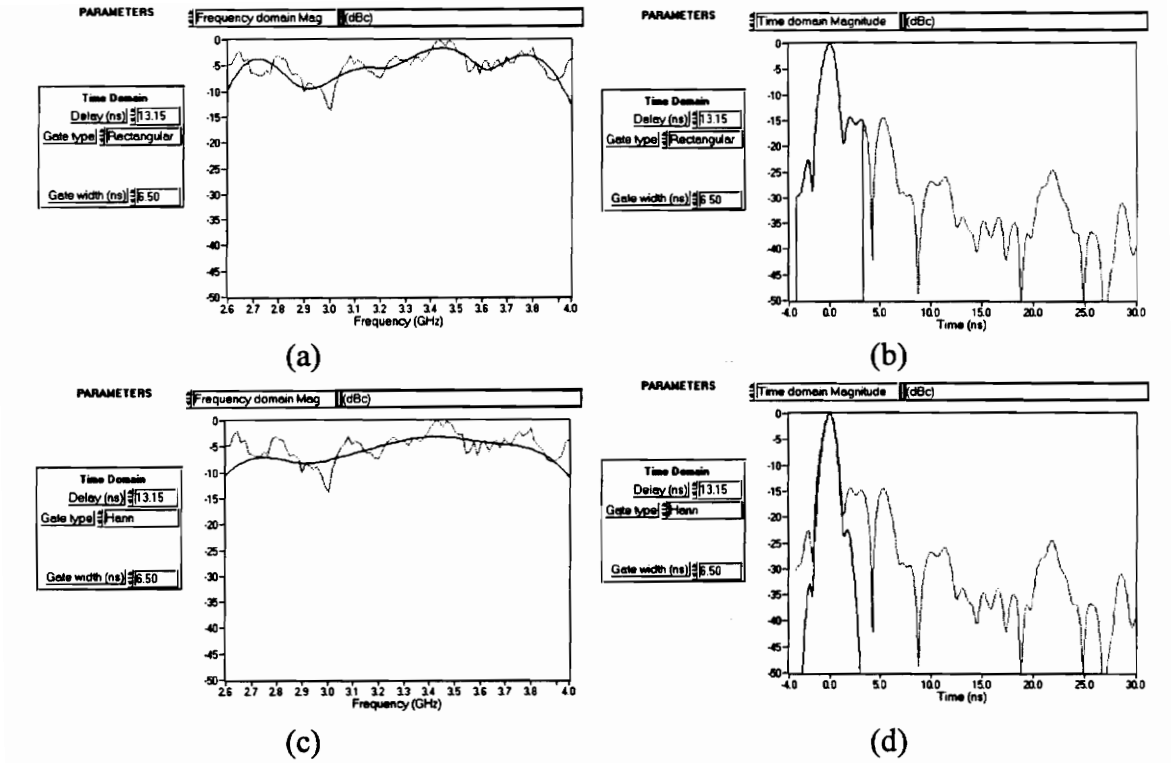

Figure 4-8: Illustration of the effect of time-domain taper. Dotted = ungated, solid = gated. (a) Frequency domain effect of rectangular time gate. (b) Rectangular time gate. (c) Frequency domain effect of Hann-tapered time gate. (d) Hann-tapered time gate.

The user can change the gate position, width, and taper while watching the results in either the frequency domain or time domain. Because the gating procedure relies on FFTs, which are quite efficient, it is practical to recompute new results every time a parameter is changed; the delay due to the recomputation is only about one second.

Many of the principles of Chap. 3 are apparent when using *GATELAB*. The dotted lines of Fig. 4-8 shows some typical-looking ungated frequency-domain and time domain data. A peak occurs in the time domain around time t=0, and then the envelope appears to show interference effects that arrive later. A 6.5 ns wide gate is positioned over the time-domain peak, in order to reject these apparent interference components. All the solid curves in Fig. 4-8 indicate gated data. Fig. 4-8 (a) and (b) correspond to an

untapered gate, and Fig. 4-8 (c) and (d) show the result of a gate shaped with a Hann taper. The effect is clear in the time domain--the gated, untapered time-domain data drop abruptly to zero outside the gate width, whereas the tapered time-domain data diminishes gradually at the edges of the gate. In the frequency domain, the untapered gate causes frequency-domain smoothing for which many residual "bumps" remain, whereas the tapered gate yields a smoother spectrum because the time-domain tapering causes the frequency-domain smoothing function to broaden, and hence smooth out the spectral variations more.

We will end this discussion by looking at some of the source code behind *GATELAB*. The *GATELAB* source code comprises about 17 modules, but the basic computation occurs in three modules shown in Fig. 4-9, 4-10, and 4-11. Fig. 4-9 shows how the Gated and Excitation data are computed from the S21 measurements. The inputs to this module are on the left side of the diagram; they are the parameters chosen by the user (gate size, etc.), and the currently-selected row of **S** (i.e., frequency-domain data at a fixed azimuth). First, delay compensation is applied to the frequency-domain data; this is useful when one wants to overlay the synthetic input and output waveforms, since the widening of the synthetic output is roughly related to the bandwidth of the antenna. The result appears as the frequency-domain Ungated curve. Next the user's chosen frequency-domain taper is computed and stored, to be shown as the frequency-domain Excitation curve. The taper is Fourier-transformed into the time domain and appears as the time-domain Excitation curve. The taper is then applied to a copy of the transmission coefficient data. The tapered transmission coefficient data is transformed into the time domain and appears as the time-domain Ungated curve. Finally, the time-domain samples of the gate function are computed (provided that gating is enabled), and the samples are multiplied with both the FD-untapered TD data and the FD-tapered TD data; the latter appears as the time-domain Gated curve. The former is transformed back to the frequency domain and shows up as the frequency-domain Gated curve. This module also performs some error checking. Note that the envelopes of the time-domain signals are in general

complex, and so only the complex modulus of the Ungated and Gated signals are actually shown.

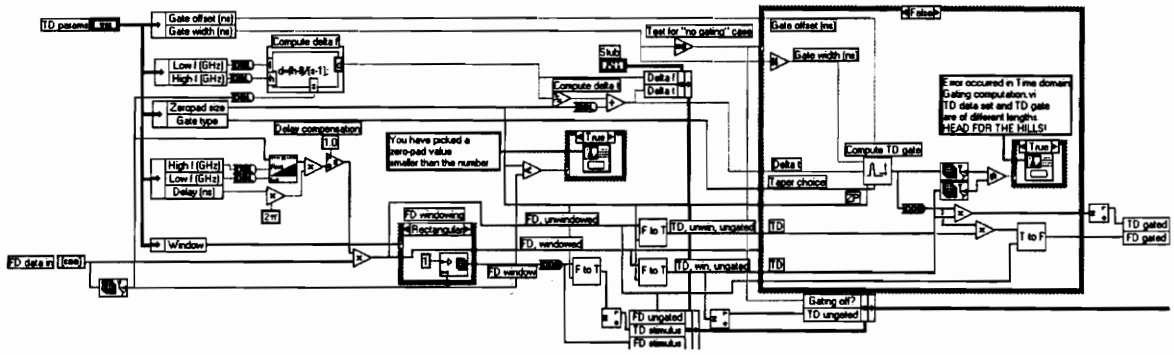

Figure 4-9: LabVIEW diagram for "Time Domain Gating computation.vi"

The transformations between the time and frequency domains are shown in Fig. 4-10 and 4-11. Fig. 4-10 shows the transformation from the frequency domain to the time domain. At one time *GATELAB* worked under the assumption that the envelope of the synthetic output was real, and so in order to display only the real part of the output envelope, all the phasors were rotated in the Argand plane so that the center frequency phasor was purely real. (This is apparently the approach taken by the HP 8409 network analyzer, although in retrospect it does not appear very sensible.) That explains some of the extra complexity in these diagrams. The remaining structures in the diagrams are required in order to put the phasor data in their proper locations, as illustrated in (3.37), in Fig. 4-10, and to extract the phasor data from the proper locations, as illustrated in (3.52), in Fig. 4-11.

To conclude, the software gating principles presented in Chap. 3 were put to use as part of a comprehensive antenna range upgrade begun and completed in 1995. This upgrade consisted of a small amount of hardware acquisition, development, and integration, and a large amount of software development--specifically, the development of a user friendly program used to run the antenna range. This program, *Range Runner*, developed by the author, performs data acquisition, instrument interfacing, and time-domain gating,

all in order to comprehensively address the demands placed upon the Virginia Tech roof-top antenna range.

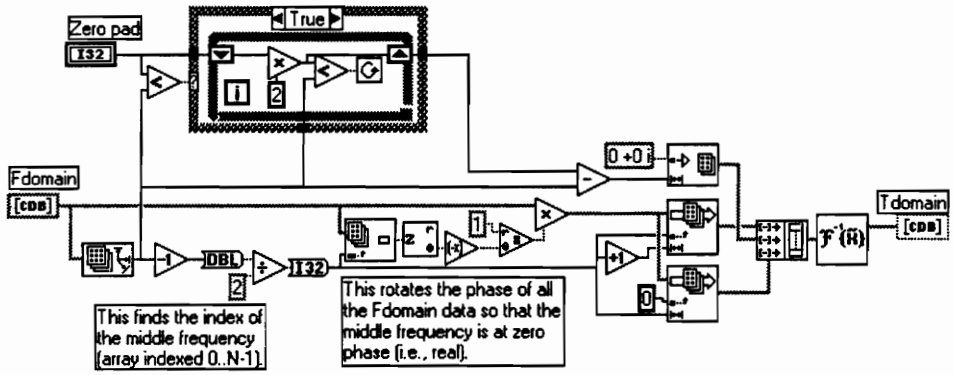

Figure 4-10: LabVIEW diagram for "Frequency domain to time domain.vi"

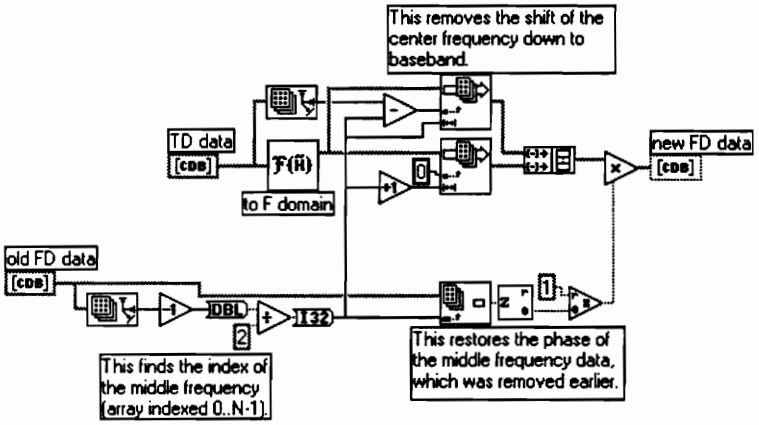

Figure 4-11: LabVIEW diagram for "Time domain to frequency domain.vi"

CHAPTER 5

VALIDATION OF THE SOFTWARE GATING SYSTEM

5.1 Overview Of The Validation Process

The validation of the upgraded system, including hardware and software, consisted of numerous tests of the system components (such as the user interface, interface to the equipment, the implementation of the equations of Chap. 3), during and after the system development. Validation of the gating system consisted largely of measuring patterns of antennas with relatively well-known patterns, and comparing the measurements to the expected results both before and after gating.

The discussion in the previous chapters has presented the numerous sources of error in the final result. To summarize, they are:

- Range multipath
- Receiver dynamic range
- Thermal drift errors
- Receiver flatness vs. frequency (same for phase reference channel, source antenna, etc.)
- Gating-algorithm-induced errors (edge effect and flattening effect)
- Gating-implementation-induced error (aliasing, finite word length)

A system that only measures CW patterns and does not employ gating will yield results subject to only the first two error results; however, the first one, range multipath, is

significant. Proper application of gating should reduce the multipath effects but not the remaining four. The data collection involved in the gating system can be quite lengthy, which can lead to measurement errors due to thermal drift. As was demonstrated in Sec. 3.5.3, the gating system requires wideband pattern data; therefore frequency-dependent sources of error enter into the results. The smoothing function introduces error, and finally the implementation of the gating can also introduce error. Additionally, as was discussed in Sec. 3.5, in circumstances such as a narrowband AUT the gating procedure is incapable of reducing multipath effects. Therefore the objective of the system validation is to demonstrate an improvement, not a total solution, to the range multipath problem.

5.2 Test Results

Both low-gain (wide beam) and high-gain (narrow beam) antennas were included in the test phase. Low-gain antennas have wide beamwidth and hence are good test subjects, since they pick up multipath from all directions. First we consider in detail an open-ended waveguide, an antenna with low gain and a bandwidth of about 50%. The second is a conical dipole antenna, with properties similar to the open-ended waveguide. Finally we look at results from a pyramidal horn, which has fairly high gain.

5.2.1 Open-Ended Waveguide

Open-ended waveguides are widely used as microwave radiators. They are used as probes for near-field antenna measurements and as a low-gain antenna standards. For this test, a section of open-ended WR-284 waveguide on a ground plane was used as the AUT. See Fig. 5-1. This type of waveguide has a cutoff frequency of about 2.07 GHz, and a recommended operating range of 2.60 to 3.95 GHz. E-plane and H-plane patterns were measured in the forward hemisphere, in 2° azimuth steps, from 2.08 to 4.00 GHz in 80 MHz frequency steps. A 5 ns Hann gate was applied to the E-plane time-domain data, and a 10 ns Hann gate was applied to the H-plane time-domain data.

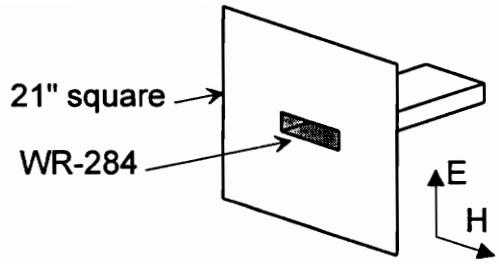

Figure 5-1: The AUT is an open-ended, probe-fed waveguide section radiating through a slot in a 21" x 21" metal plate. The slot and the waveguide are the same size.

An open-ended waveguide operating in the TE_{10} mode has a pattern that can be predicted by applying aperture antenna theory; according to [Stu81], the approximate patterns for the infinite ground plane case are:

$$F_E(\theta, f) = \frac{\sin\left(\frac{\pi f}{c} b \sin\theta\right)}{\frac{\pi f}{c} b \sin\theta}, \quad |\theta| < \frac{\pi}{2} \tag{5.1}$$

$$F_H(\theta, f) = \cos\theta \cdot \frac{\cos\left(\frac{\pi f}{c} a \sin\theta\right)}{1 - \left(\frac{2f}{c} a \sin\theta\right)^2}, \quad |\theta| < \frac{\pi}{2} \tag{5.2}$$

where a is the width of the broad wall of the waveguide, b is the width of the narrow wall, c is the speed of light, and θ is the azimuth measured away from boresight. These are normalized patterns, and do not include antenna gain.

Figs. 5-2, 5-3, and 5-4 show a comparison of the measured E-plane patterns, before (U) and after gating (G), to (5.1), at three different frequencies. The vertical scale is in "dBx," which denotes decibels relative to an unknown reference defined inside the SA 1780 receiver (cf. Sec. 3.1). For ease of comparison, the theoretical pattern (T) intersects the gated pattern at boresight; these graphs are not meant to show that the absolute gain measurement is more accurate after gating, although we expect that result. Significant

ripple is evident at all three frequencies both before and after gating; nevertheless, the agreement with the theoretical pattern apparently improves after gating. Note that the ripples are not necessarily artifacts of range multipath; they could be caused or aggravated by the finite ground plane, which is about 5 wavelengths per side at the lower frequency range. In observing the effect of gating, one notices that the gross departures from theory are mitigated by the gating; in Fig 5-2, there is a large dip around -25° which the gating removes. In Fig. 5-3, a large dip around 10° is removed by the gating, and in Fig 5-4 the gated result seems to indicate that the entire sector from -20° to 90° is erroneously shifted upwards, although this probably due to the edge effect (Sec. 3.5.3), since the ungated and gated reference levels are only comparable from 2.3 to 3.8 GHz, approximately. In fact the dip from -90° to -20° in the ungated pattern is probably due to multipath distortion.

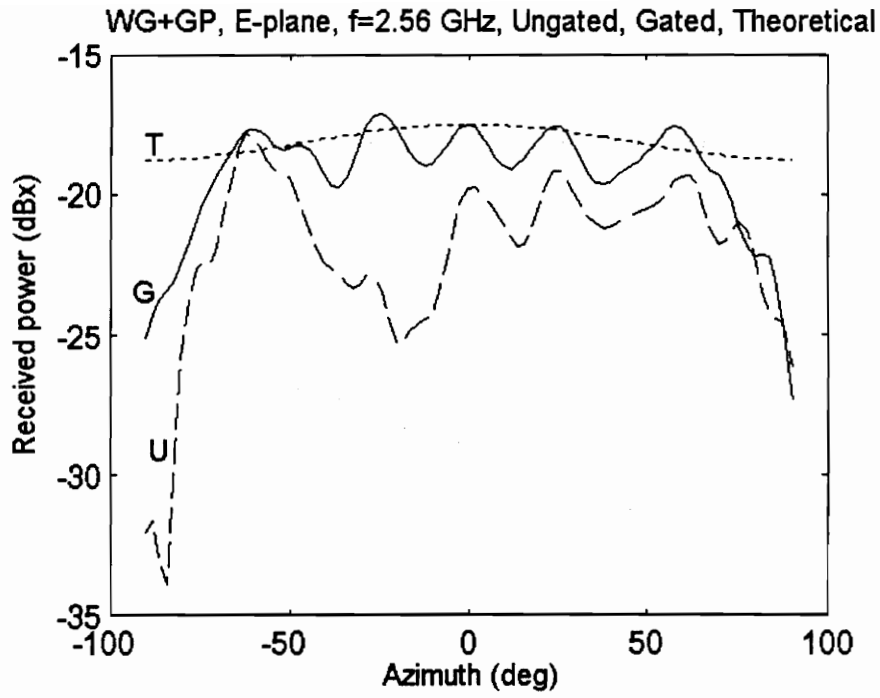

Figure 5-2: Waveguide E-plane patterns at 2.56 GHz. U=Ungated pattern, G=Gated pattern (using a 5 ns wide Hann gate), T=Theoretical pattern (eq. 5.1).

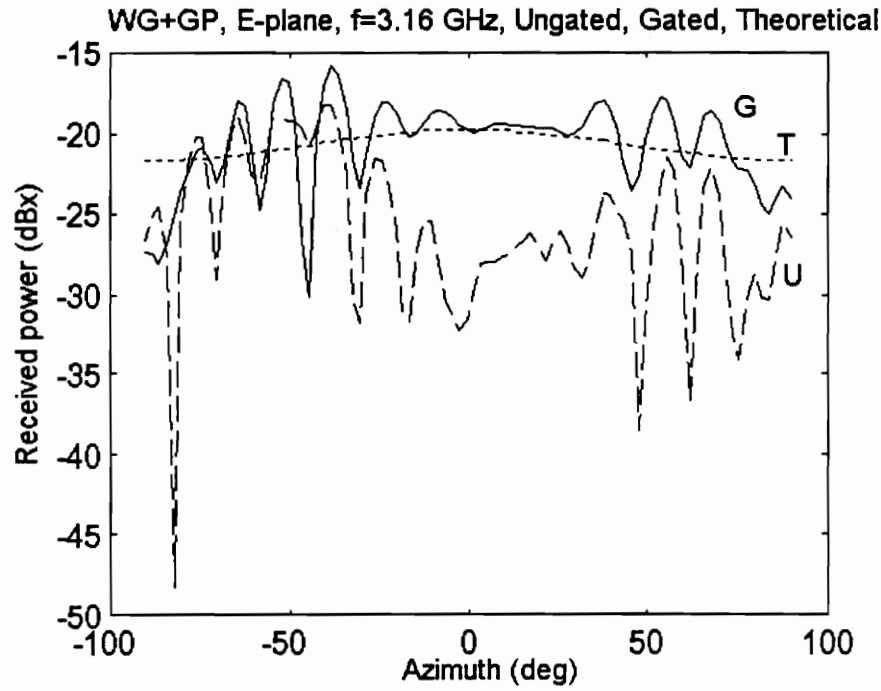

Figure 5-3: Waveguide E-plane patterns at 3.16 GHz. U=Ungated pattern, G=Gated pattern (using a 5 ns wide Hann gate), T=Theoretical pattern (eq. 5.1).

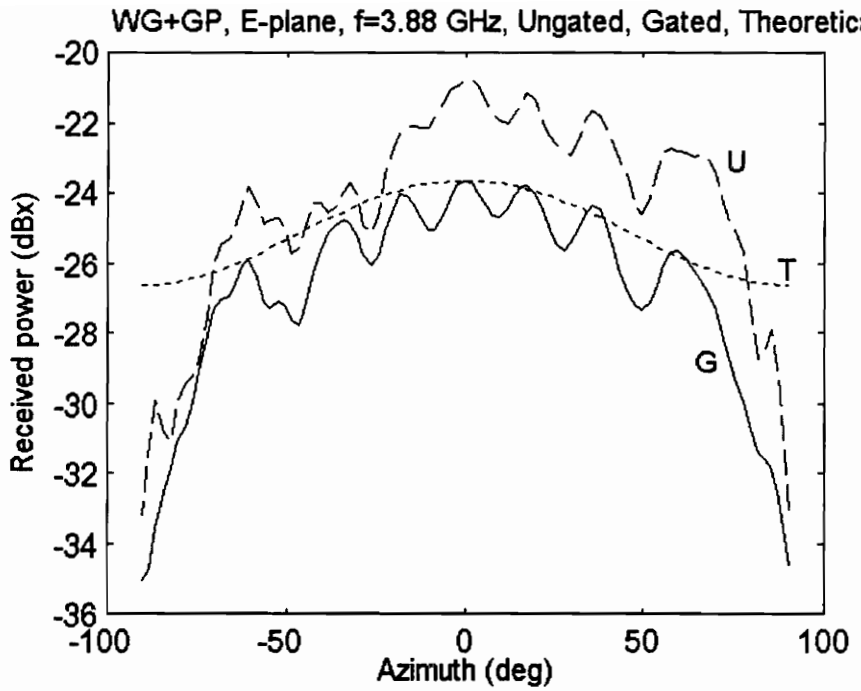

Figure 5-4: Waveguide E-plane patterns at 3.88 GHz. U=Ungated pattern, G=Gated pattern (using a 5 ns wide Hann gate), T=Theoretical pattern (eq. 5.1).

Fig. 5-5 shows the measured pattern data over the measured frequency range. The ungated pattern is shown in the surface in Fig. 5-5 (a). The surface varies rapidly versus frequency and azimuth, which is not the expected behavior of this antenna. After the gating, the surface is much smoother, as seen in Fig. 5-5 (b). Note that the surface is smoother with respect to azimuth angle as a result of the gating, and azimuthal symmetry has increased, even though (as discussed in Chap. 1) the gating procedure is "blind" with respect to azimuth. This helps confirm that Fig. 5-5 (b) is closer to the true antenna pattern. The edge effect is probably evident in Fig. 5-5 (b); clearly the surface is attenuated towards the terminal frequencies, although the waveguide itself causes part of this attenuation. Features not otherwise visible appear in the gated data; for example, the dip in the response at 3.6 GHz was obscure in the ungated surface, but is much more obvious in the gated surface. Note that symmetry is not perfect; this suggests that there are still

problems with range multipath. The gating procedure is clearly not a cure-all, but it does seem to improve the pattern estimates.

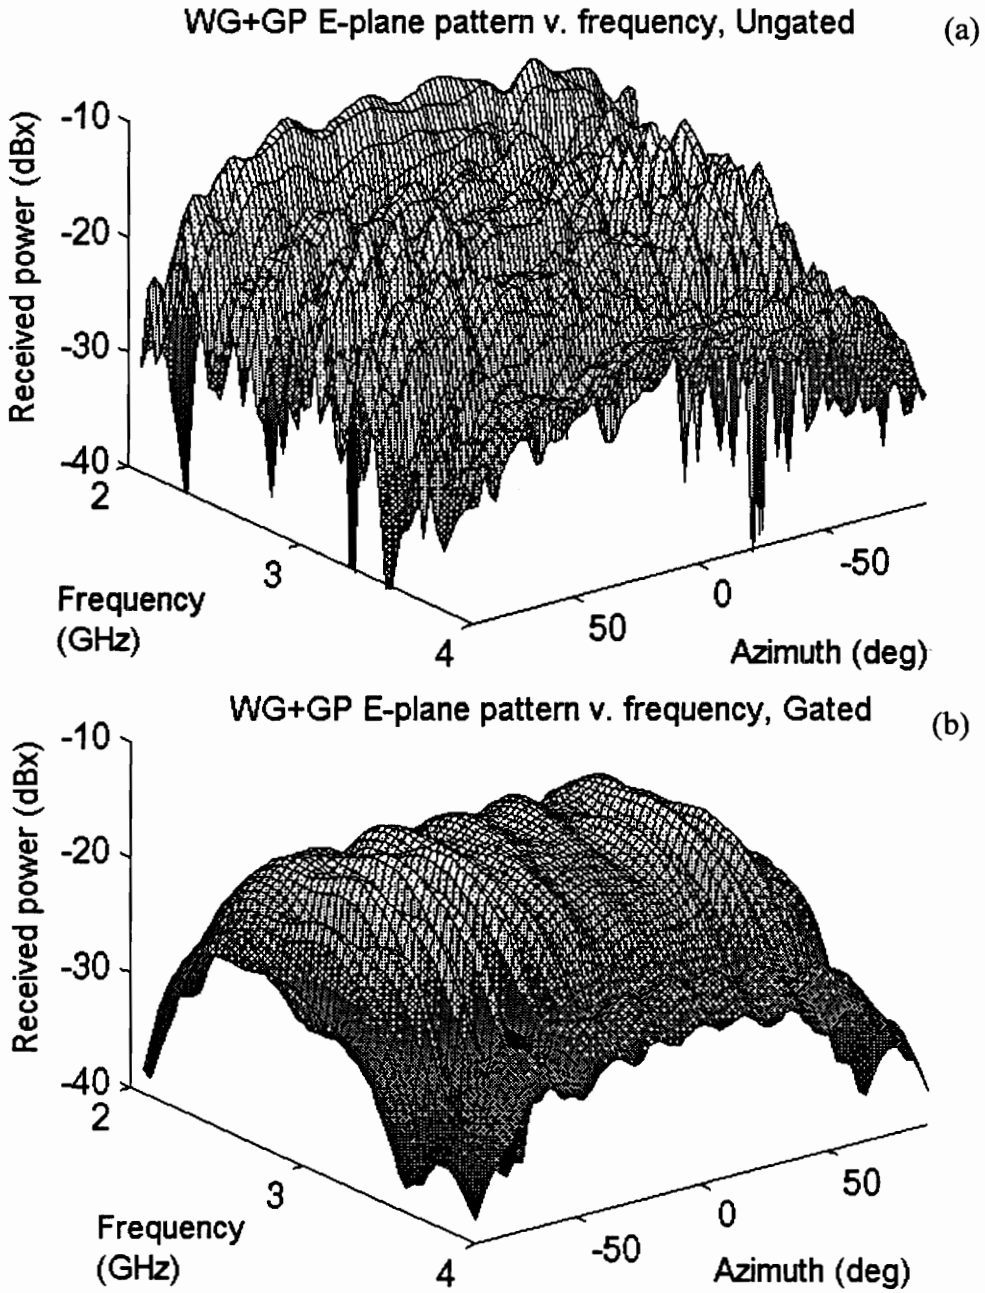

Figure 5-5: Waveguide E-plane patterns v. frequency. (a) Ungated. (b) Gated, using a 5 ns wide Hann gate.

Results in the H-plane of the antenna are shown in Figs. 5-6, 5-7, and 5-8. Again, the gated results show better fidelity to the theoretical patterns than the ungated patterns, which indicates that the gating is beneficial. Nevertheless, especially in Fig. 5-8, there are still problems with the gated patterns, as seen by the ripple and asymmetry in the final result.

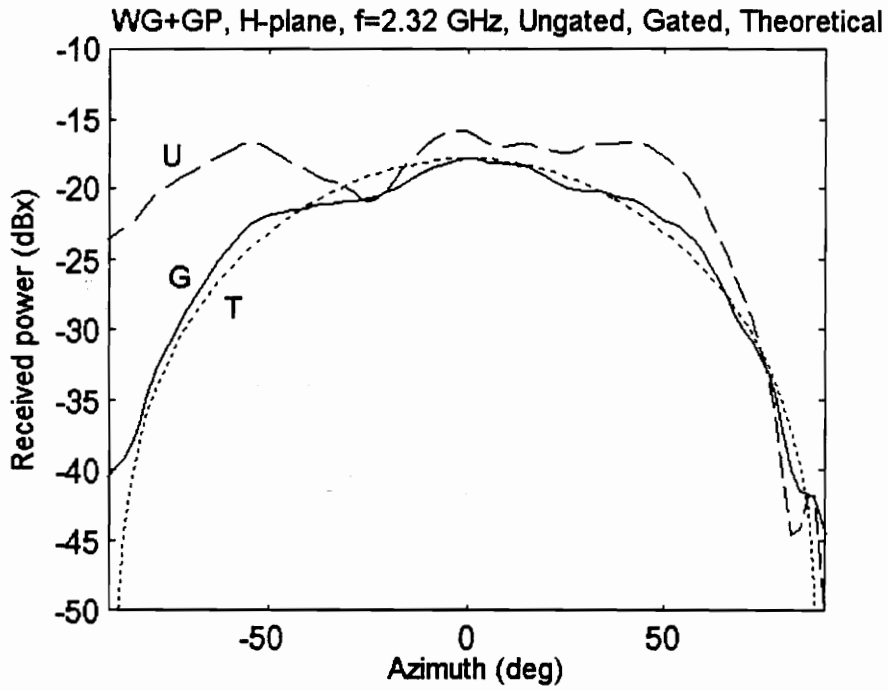

Figure 5-6: Waveguide H-plane patterns at 2.32 GHz. U=Ungated pattern, G=Gated pattern (using a 10 ns wide Hann gate), T=Theoretical pattern (eq. 5.2).

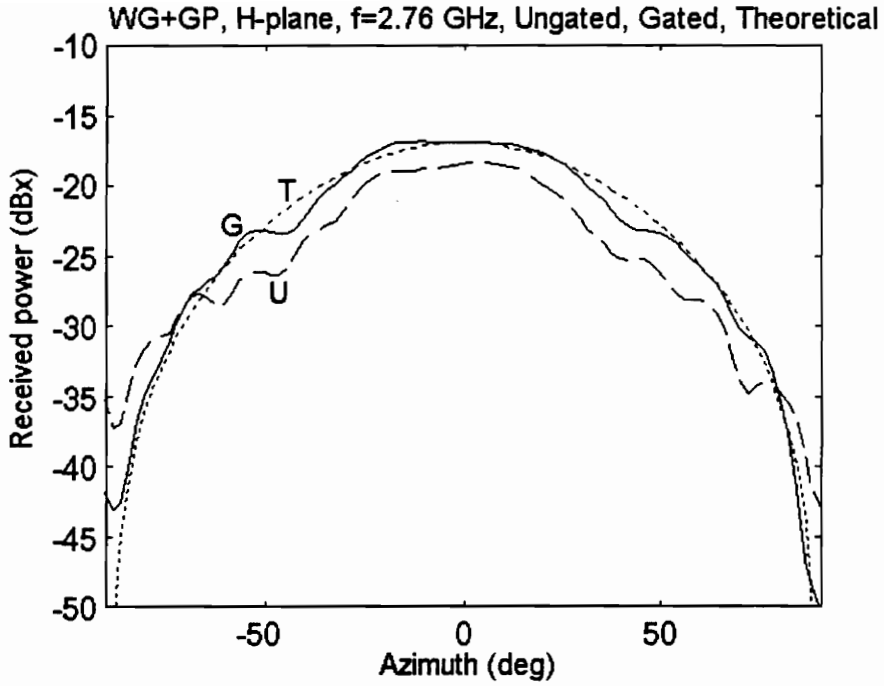

Figure 5-7: Waveguide H-plane patterns at 2.76 GHz. U=Ungated pattern, G=Gated pattern (using a 10 ns wide Hann gate), T=Theoretical pattern (eq. 5.2).

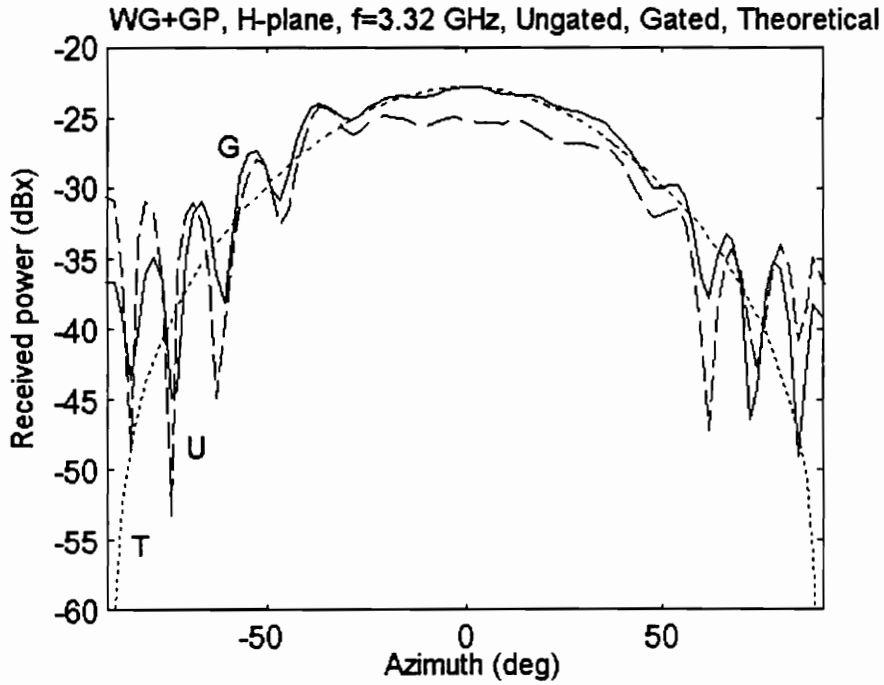

Figure 5-8: Waveguide H-plane patterns at 3.32 GHz. U=Ungated pattern, G=Gated pattern (using a 10 ns wide Hann gate), T=Theoretical pattern (eq. 5.2).

Fig. 5-9 shows pattern v. frequency surfaces for the ungated and gated H-plane data. Conclusions similar to those about the E-plane pattern surfaces can be made. The fluctuations in level with respect to frequency apparent in Fig. 5-9 (a) are not expected of this antenna; this indicates that gating (that is, frequency-domain smoothing) is justified. Again, features that were obscured by the interference are visible in Fig. 5-9 (b), such as the dips in the frequency response at 3.2 and 3.6 GHz. The smoothness of the gated surface in the azimuth direction is a property hidden in the ungated surface; that is, the gating algorithm does no filtering in the azimuth direction, but only in the frequency domain, and when the frequency-domain ripple is smoothed, the smoothness in the azimuth direction can be seen.

This waveguide example serves as a good validation of the gating system. Clearly the gating performed on the E-plane and H-plane measurements has improved the fidelity of the pattern measurements compared to the ungated data, using the simple theoretical pattern as an ideal. The remaining examples show a few other interesting facets of the gating system that are not covered in this example.

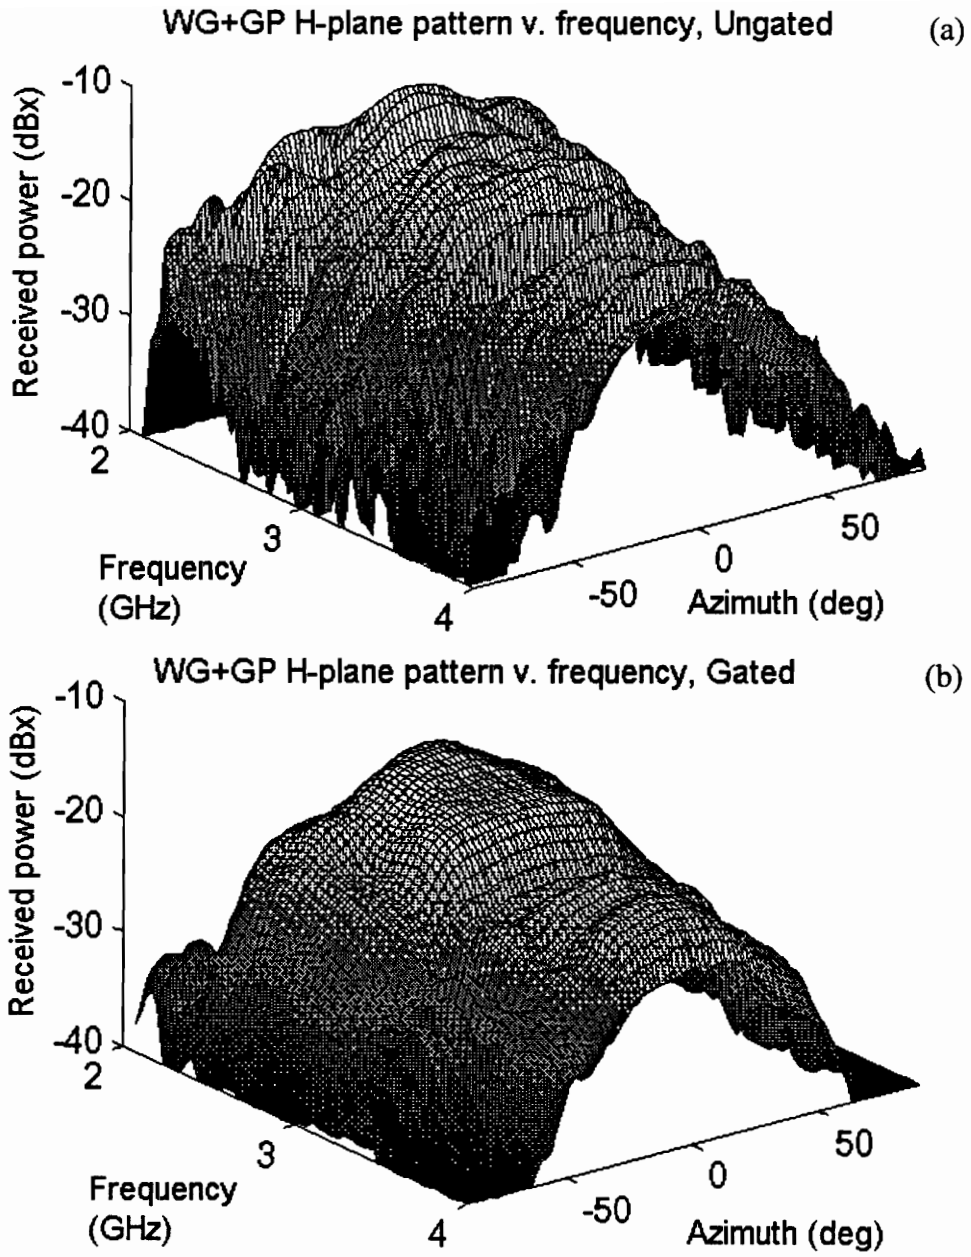

Figure 5-9: Waveguide H-plane patterns v. frequency. (a) Ungated. (b) Gated.

5.2.2 Tests With A Conical Dipole

A dipole antenna made of two cones, such as the one illustrated in Fig. 5-10, has better bandwidth than a similar-length wire dipole. The pattern of the conical dipole

shown was measured from 1 to 2 GHz in 15.625 MHz steps. The problems that occur measuring this antenna's patterns are similar to those of the last antenna, but it is worth our attention because the ungated pattern shows a phenomenon called *null filling* well. According to [Stu81] a conical dipole has a pattern very similar to a wire dipole, and a wire dipole of course has distinct nulls off its ends. One common manifestation of range multipath is that, when measuring a dipole pattern, the nulls in the measured pattern are of finite depth. Fig. 5-11 shows the gated and ungated E-plane pattern of the antenna at 1.2 GHz, and clearly the gated pattern has much deeper nulls. Thus one quantifiable benefit of the gating system is that the measurable null depth increases by 10 dB or more.

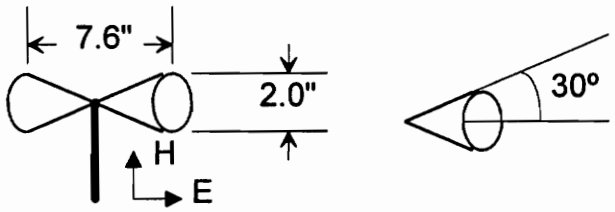

Figure 5-10: Conical dipole test article. Cone angle is 30°

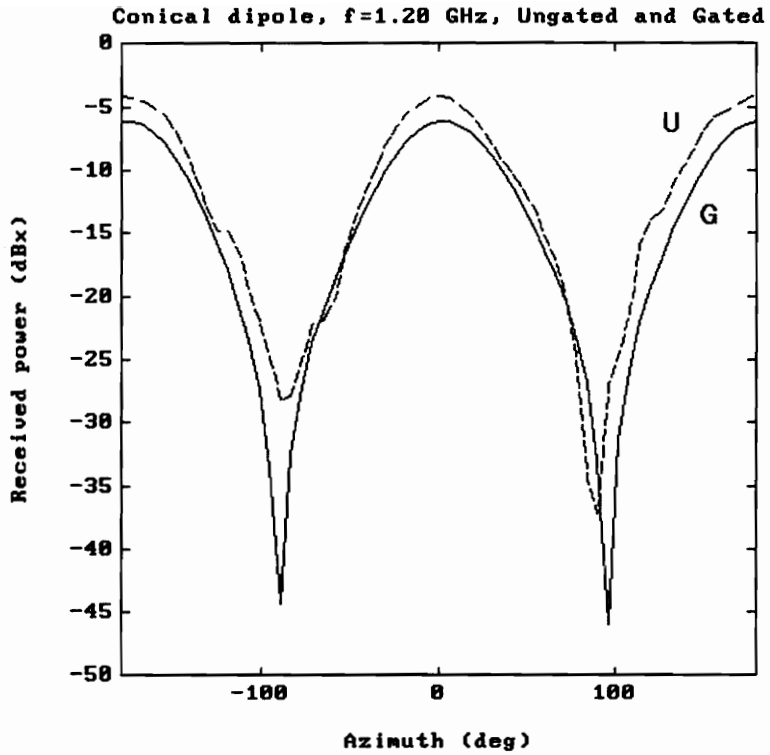

Figure 5-11: Conical dipole patterns at 1.20 GHz, gated with a 15 ns Blackman gate.

5.2.3 Pyramidal Horn

The surfaces shown in Fig. 5-5 and 5-9 are functions azimuth - frequency domain. This is the natural format of the *Range Runner* output file. Of course, the discrete Fourier transform can be applied to the frequency-domain data (with or without tapering) and hence the pattern data can be shown as a surface in the azimuth - time domain. We must briefly discuss the interpretation of this presentation now.

If the synthetic input to the channel is narrow compared to the bandwidth of the channel, then this surface approximates the time-domain impulse response of the channel versus azimuth. An assumption underlying the ray model which we have been using since Chap. 1 is that sources of interference on the range will be localized, and hence have a specific angle of arrival (an azimuth coordinate) and a specific delay (a time

coordinate). Of course, the synthetic pulses excite the channel, not just the antenna; that is, they excite the interference locations on the range. The range and antenna are in effect cascaded; therefore the impulse response of the cascade will equal the convolution of the impulse responses of the channel and test antenna. This is illustrated in Fig. 5-12.

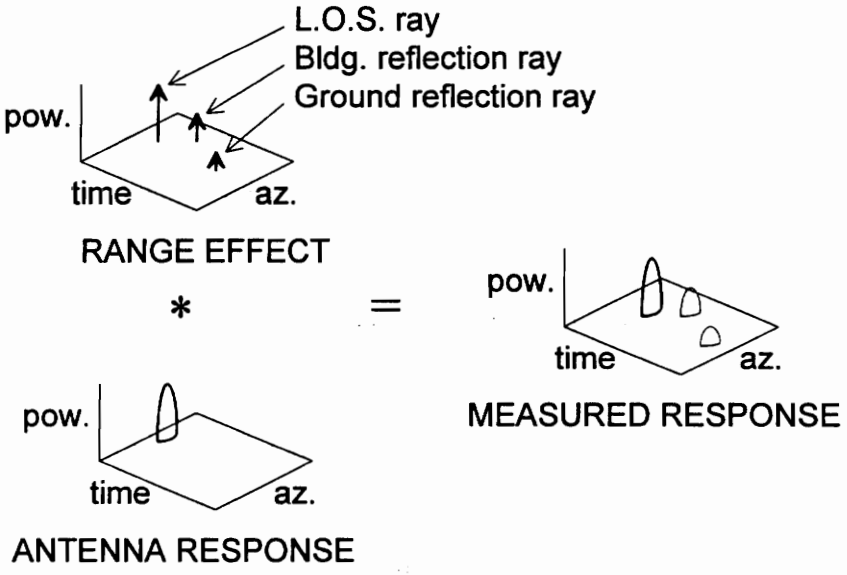

Figure 5-12: In the time - azimuth domain, the measured channel response is a convolution of the antenna response and the range response.

Hence, if an appropriate antenna is measured, its azimuth - time domain surface will point to sources of interference. Fig. 5-13 shows the time - azimuth domain channel response when a Seavey SGA-20 pyramidal horn antenna is the test antenna. Since this antenna is relatively directive, its response is fairly narrow in the azimuth domain; furthermore, since it is wideband, its response is narrow in the time domain. So, this antenna should be fairly successful in revealing the origin of rays. Some of the peaks in Fig. 5-13 are therefore labeled with hypotheses of the source of the exciting ray: ground bounce interference is usually going to occur at zero azimuth, whereas reflections off the penthouse are going to have a significantly different angle of arrival. Both will be delayed relative to the line-of-sight component. Peaks on this surface are found near the

expected locations for these antenna responses. Thus this plot helps validate the gating system because it seems to confirm the hypothesized sources of ray interference.

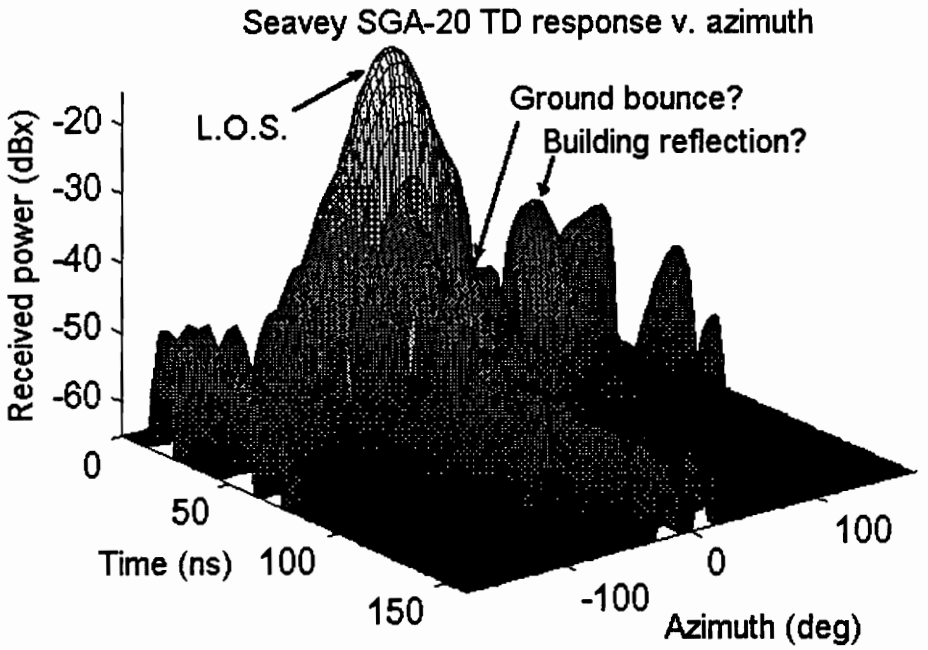

Figure 5-13: Pyramidal horn response v. time and azimuth.

To sum up, this chapter has presented results from when the gating system is applied to pattern measurements. We have seen that although gating does not solve all the problems in the antenna range, it does improve the fidelity of the measured patterns, and provides insight into the behavior of the antenna range.

CHAPTER 6

RECOMMENDATIONS AND CONCLUSIONS

This chapter lists a few directions in which the antenna range work could be continued, and concludes the thesis with a summary of the improvements and limitations of the gating method.

6.1 Conclusions

This thesis has presented the principles of a software gating system, discussed its implementation and validation, and examined some of the errors that occur as a result of gating. Frequency-domain measurements (Sec. 3.1) can be combined with the IDFT (Sec. 3.4) to yield channel response to a synthetic input. The channel response can then be gated and transformed via the DFT to yield frequency-domain results (Sec. 3.5). These equations were built into a software gating system (Sec. 4.3.2) which was validated for a number of antennas (Sec. 5.2). This system is very similar to network-analyzer applications (Sec. 2.4). The potential error mechanisms of gating are at least threefold (Sec. 3.5.3), although only one of these mechanisms has been previously discussed (Sec. 2.5).

6.1.1 Summary Of Improvements

- The Virginia Tech antenna range was upgraded to perform vector antenna measurements via a highly automated system that can span a range of frequencies and azimuths.
- The *Range Runner* program is a user-friendly data collection program.

- The channel response can be seen in the time domain; the user has a choice of synthetic pulse shapes; the gate size, location, and shape can be changed easily; the user can see the effect of the gate in both the frequency domain and time domain easily.
- Pattern fidelity has improved. Qualitatively, gated patterns agree better with the patterns predicted by theory than do the ungated patterns. Quantitatively, measurable dipole null depth has increased from about 20 dB to about 40 dB.
- Time domain data reveal insight into the behavior of the antenna range and suggest new ways that the range effects may be mitigated.

6.1.2 Limitations Of The Software Gating Method

- Narrowband AUTs have time-domain responses that are too slow, relative to the times of arrival of the multipath components, to permit interference reduction.
- Frequency-domain smoothing can obscure or distort AUT response.
- A large amount of time is required to fill the 2D array with channel data. This is inconvenient, and permits thermal drift errors to accumulate.

6.2 Future Work

There are many features that could be added to the antenna measurement system to improve its usefulness or ease of use. The system hardware can be more completely automated, the system software can be enhanced to make certain measurements more convenient, the measurement results can be more fully utilized to reveal information about the range behavior, and perhaps to cancel out some of the unwanted range behavior.

6.2.1 Increased System Automation

The antenna range requires a large amount of manual setup effort, regardless of the measurement itself. Some of the items that presently require manual setup impose a limit to what the user can achieve with the system; but, with some additional hardware, these items could be automated and the system potential could be more completely used.

The upgraded system of Fig. 4-1 (b) uses one of two synthesizers to excite the channel; one covers 100 kHz to 3.2 GHz, the other 2.0 GHz to 26.0 GHz. However, the user must manually connect the desired synthesizer to the test channel with a piece of coax. Thus it is presently impossible for the user to make swept-frequency measurements in a band that spans the overlap between synthesizers (for example, a sweep from 1.5 to 3.5 GHz is impossible). It would be nice to automate this switchover, perhaps by the use of coaxial switches under computer control; this upgrade would not be too expensive and it would increase the usefulness of the antenna range.

The SA 1780 receiver uses a special Low Frequency Converter (LFC) when it is measuring patterns from 100 kHz to 1 GHz. When working in this frequency range, the LFC replaces the usual harmonic mixer; the user must manually connect the channel output to the proper mixer. Therefore it is impossible to make swept-frequency measurement spanning the 1 GHz boundary. Again, it would be good to automate this switchover, perhaps with coaxial switches under computer control.

On the source antenna tower there is a manually controlled polarization positioner for the test antenna. With some effort this positioner could be put under computer control, and then the system could automatically measure the axial ratio of the antenna polarization. This would be especially useful when measuring circularly-polarized antennas such as helices.

6.2.2 Antenna Gain Measurements

Antenna gain is measured on the Virginia Tech antenna range using the substitution method; the reference level of the test antenna pattern is compared to the reference level of a standard gain antenna. The gain standard must be measured anew each time this comparison is performed, since the channel output level is quite variable in the long term; it is not only affected by the channel input amplitude and path loss, but also by thermal effects and aging, the choice of cables used to set up the experiment, etc. So, to measure antenna gain for a particular AUT always requires a pattern measurement on the AUT plus a boresight transmission coefficient measurement on a standard gain antenna. Then the results are compared, which yields an estimate of the gain of the AUT. *Range Runner* could be enhanced to automate this procedure. (Note that gating will benefit not only the AUT measurement but also, to a lesser extent, the boresight measurement on the gain standard.)

6.2.3 Visualization / Characterization Of Range Phenomena

As discussed in Sec. 5.2.3, when the pattern data are transformed into the time - azimuth domain and plotted as a surface (e.g., see Fig. 5-13), localized sources of interference are apparent. This application of the software has not been fully explored; it could be useful for visualizing the sources of interference on the range, for suggesting locations for fences or absorber, and for validating other applications of range-quieting techniques. Because the test tower can only vary azimuth, this approach will only reveal the azimuth angle of arrival of interference. If the test tower were upgraded to an Az-El pedestal, then the full angle of arrival (that is, θ and ϕ) of ray interference could be estimated. This leads directly to the next idea: cancellation, via an equalizer.

6.2.4 Development Of A Range Equalizer

As established in Sec. 5.2.3, in the time - azimuth domain the AUT response and the range response will be combined by convolution, as illustrated in Fig. 5-12. Again we use the fact that the Fourier transform relates convolution to multiplication. If the time - azimuth domain data are 2D-Fourier transformed to the frequency - aperture domain, the range effect can be expressed as a factor, as shown in Fig. 6-1.

If the range effect factor were isolated, its inverse could be used as an equalizer. Then the AUT measurements could be cleaned up systematically, without the ad hoc step of choosing a gate size, simply by performing the appropriate transformation and multiplying the results by an equalizer. This suggests a potential solution, but one must be careful not to be misled by Fig. 6-1. This approach is not feasible on the rooftop antenna range as it is now configured, because in general the range effect factor cannot be isolated without both azimuth and elevation pattern data. If one currently tries to determine the range effect factor by using a reference antenna, and then to equalize the measurements from a test antenna, the difference in elevation antenna pattern between the antennas will cause the range effect to be different for the two antennas, and the equalizer applied to the AUT data will not be appropriate.

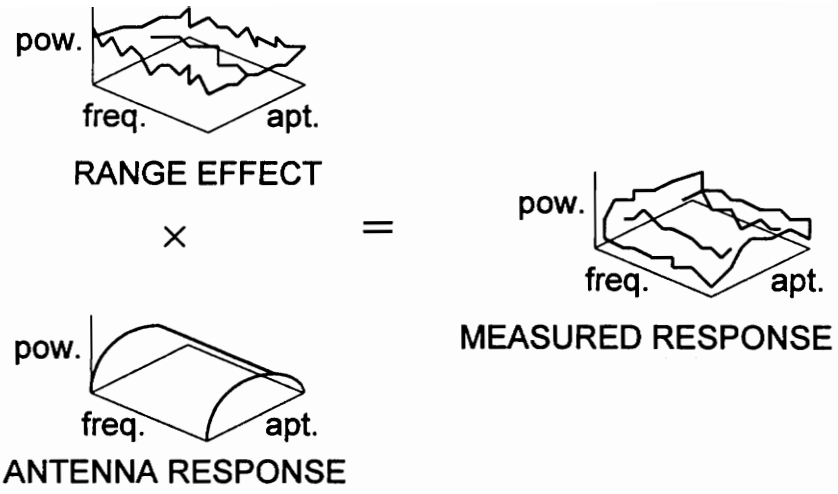

Figure 6-1: Range effect in the frequency - aperture domain. Each of these is the 2D Fourier transform of the corresponding surface in Fig. 6-1.

6.2.5 Application Of Spread-Spectrum Interference Cancellation Techniques

Another approach to interference reduction that has been explored is the use of spread-spectrum (SS) interference cancellation techniques. The approach taken by Zeger and Abrams [Zeg95] is to excite the channel with a SS signal, attach a SS receiver to the AUT, and use the output to drive an LMS estimation loop controlling the parameters of a model of the channel. The model parameters reveal the gain of the line-of-sight ray, and hence the antenna pattern. The signal that excites the channel is, of course, band-limited, and probably periodic. If this is so, then perhaps the channel could be characterized by **S**, and a synthetic SS-like signal found in D, the domain of the channel transformation (c.f. eq. 3.19). Then, the interference cancellation could be done in software, which provides the many advantages already discussed.

APPENDIX A: GLOSSARY

This appendix presents definitions of terms used in this thesis. Some of the terms have a meaning that is more restricted than that of general usage.

Antenna gain

The ratio of power available from the antenna (at a particular angle, with plane-wave excitation at a particular frequency), compared to the power available from an (idealized) isotropic receiving antenna. See [Stu81] for a comprehensive definition.

Antenna pattern

The power received by an antenna as a function of spatial orientation. Patterns are usually normalized to have a maximum at 0 dB.

Antenna range

A facility used for measuring antenna patterns. See also Far-field range and Compact range.

Antenna under test (AUT)

The antenna for which the pattern is being measured. The AUT is configured as the receive antenna, and mounted on a positioner that controls antenna azimuth.

Aperture (effective aperture)

Effective collecting area of an antenna. An antenna's effective aperture depends on polarization, orientation, and frequency. See [Stu81] for a comprehensive definition.

Boresight

A line perpendicular to the azimuth axis of the antenna positioner and parallel to the Poynting vector of the (idealized) EM plane wave illuminating the test zone.

Channel

Usually this refers to the two-port network between the source antenna and the AUT, comprising the antenna range, along with all its sources of reflection and scattering. But see also Reference channel. If there is any ambiguity about which channel is meant, the one involving the AUT will be called the test channel. Both channels are assumed to be linear systems.

Compact range

An indoor facility for measuring antenna patterns, consisting of a source antenna (usually a large reflector fed with a horn) and a AUT. The source antenna is designed to illuminate the AUT with a good approximation of a plane wave, even though the distance between the two antennas is much less than the far-field distance of the source antenna.

Cut

A pattern measurement in which the antenna orientation is only varied about a single axis.

CW (continuous-wave) excitation

A sinusoidal input to a channel.

CW response

Measurements showing the response of the channel to CW excitation, usually at a series of evenly spaced frequency steps.

Far-field range

A facility for measuring antenna patterns, consisting of a source antenna and an AUT, situated at a distance from one another in an open space. The two antennas are distant enough that both are in each other's far field.

Gate function

A function of time used to exclude portions of the pulsed response. When the gate function is near or at zero, the contemporaneous part of the pulsed response is nearly or completely excluded. The excluded portion is one interval per pulse period, with start and end times chosen to reduce the effects of interference.

Gated response

See Gating.

Gating

The partial or complete exclusion of a portion of pulsed response, by multiplication with a gate function. The product of the pulsed response and the gate function is called the gated response. Usually the purpose of gating is to reduce the effects of multipath in an antenna range and therefore improve the quality of an antenna pattern measurement.

Hardware gating

Gating implemented with RF hardware, e.g., by means of a high-speed switch. This is always associated with pulsed excitation.

Illuminate

To transmit electromagnetic energy onto an object.

Line-of-sight energy

The EM energy radiated from the source antenna into the AUT along a straight line.

Multipath

The phenomenon in which RF energy arrives at a location (such as the test zone of an antenna range) by means of a variety of paths, due to effects such as line-of-sight propagation, reflection off nearby objects, and diffraction.

Null

A point in a pattern (ideal or practical) where (ideally) the received power goes to zero, that is, negative infinite decibels. Nulls are difficult to measure; they always get "filled in," if not due to interference, then because the measurement receiver has a finite noise floor. Usually interference is the dominant cause of null filling.

Pattern

See *antenna pattern.*

Phase reference signal

In practice, the phase of a sinusoid can only be measured with respect to that of another sinusoid at the same frequency. Therefore a phase detector has two inputs: a information signal to measure, and a phase reference signal to provide a phase reference; the detector output is the phase difference between the two signals.

Positioner

A device used to rotate an antenna.

Pulse duration

The length of time containing most of the energy in a single period of the RF pulse train.

Pulse period

The period of the RF pulse train.

Pulsed excitation

An RF pulse train as the input to a channel.

Pulsed response

Measurements showing the response of the channel to an RF pulse train; although the response is periodic, typically only one pulse period is shown. Note that pulsed excitation is not required to determine the channel's pulsed response; on the contrary, the pulsed response can be synthesized from the response to appropriately-chosen CW excitation.

Reference antenna

An antenna that is used to receive the phase reference signal, when the reference channel is implemented with antennas. The reference antenna is an third antenna added to the antenna range, co-located with the AUT, but non-rotating.

Reference channel

The two-port network used to provide a phase reference signal to the receiver. Usually this is implemented with a length of coaxial cable. Alternatively, the network can be "wireless," by adding a third antenna to the range: see Reference antenna.

Reference level

A datum accompanying a normalized pattern cut equal to an absolute power measurement made at the peak of the cut; this permits different cuts to be compared to each other. It is, however, dependent upon factors independent of the AUT, such as the source antenna gain, the transmitter power, and so on.

RF pulse train

A periodic signal consisting of an amplitude-modulated sinusoid, for which most of the energy is concentrated in a single, short interval of its period. As a result, it is a signal with a characteristic center frequency and bandwidth.

Sidelobe

An undesired feature in a directional antenna pattern; specifically, a local maximum in the pattern surrounding the main beam of the pattern.

Software gating

Gating implemented in a computer; the pulsed response is multiplied by a computer-generated gate function. This is usually associated with CW excitation.

Source antenna

The antenna that provides the channel excitation for the pattern measurement, i.e., the transmitting antenna. It is stationary for most pattern measurements.

Squint

Beams or nulls are said to "squint" if they change position in the pattern as frequency changes.

Synthetic pulse

A hypothetical signal that, if it were applied to the channel, would yield the pulsed response as determined by computer processing.

Taper

A function with a finite nonzero portion, shaped at the edges to reduce the effect of the Gibbs phenomenon. (Also known as a window function.)

Tapered gate

A gate function with a smooth transition between its included and excluded intervals.

Test channel

See Channel.

Test zone

The area of an antenna range typically occupied by the AUT. For an ideal far-field or compact range, with CW excitation, the EM field in the test zone is a plane wave with uniform amplitude, phase, and Poynting vector at all points in the test zone. Often the quality of an antenna range is evaluated by sampling the field in the test zone and then plotting the amplitude variation.

Time gating

See Gating.

APPENDIX B: DERIVATION OF (3.24) FROM (3.23)

Equation (3.23), repeated below in (B.1), is the starting point for this derivation; it is a sum of N sinusoids with unity amplitude and uniform frequency spacing Δf, starting at a low frequency of f_L, and all in phase together at time t=0:

$$r(t) = \sum_{k=0}^{N-1} \cos[2\pi t(f_L + k\Delta f)] \tag{B.1}$$

$$= \sum_{k=0}^{N-1} \cos(2\pi t f_L + 2\pi t k\Delta f)$$

$$= \sum_{k=0}^{N-1} \cos(2\pi t f_L)\cos(2\pi t k\Delta f) - \sin(2\pi t f_L)\sin(2\pi t k\Delta f)$$

$$= \cos(2\pi t f_L)\sum_{k=0}^{N-1} \cos(2\pi t k\Delta f) - \sin(2\pi t f_L)\sum_{k=0}^{N-1} \sin(2\pi t k\Delta f) \tag{B.2}$$

Eq. (2.111) found in [Opp89] proves to be useful:

$$\sum_{n=-M_1}^{M_2} e^{-j\omega n} = e^{-j\omega(M_2-M_1)/2} \cdot \frac{\sin[\omega(M_1+M_2+1)/2]}{\sin(\omega/2)}. \tag{B.3}$$

We apply (B.3) to sines and cosines as follows:

$$\sum_{n=-M_1}^{M_2} \sin(\omega n) = \frac{1}{2j}\sum_{n=-M_1}^{M_2} e^{j\omega n} - e^{-j\omega n} \tag{B.4}$$

$$= \frac{1}{2j}\left\{ e^{j\omega(M_2-M_1)/2} \cdot \frac{-\sin[\omega(M_1+M_2+1)/2]}{-\sin(\omega/2)} - e^{-j\omega(M_2-M_1)/2} \cdot \frac{\sin[\omega(M_1+M_2+1)/2]}{\sin(\omega/2)} \right\}$$

$$= \sin[\omega(M_2-M_1)/2] \cdot \frac{\sin[\omega(M_1+M_2+1)/2]}{\sin(\omega/2)} \tag{B.5}$$

$$\sum_{n=-M_1}^{M_2} \cos(\omega n) = \frac{1}{2}\sum_{n=-M_1}^{M_2} e^{j\omega n} + e^{-j\omega n} \tag{B.6}$$

$$= \frac{1}{2}\left\{ e^{j\omega(M_2-M_1)/2} \cdot \frac{-\sin[\omega(M_1+M_2+1)/2]}{-\sin(\omega/2)} + e^{-j\omega(M_2-M_1)/2} \cdot \frac{\sin[\omega(M_1+M_2+1)/2]}{\sin(\omega/2)} \right\}$$

$$= \cos[\omega(M_2-M_1)/2] \cdot \frac{\sin[\omega(M_1+M_2+1)/2]}{\sin(\omega/2)} \tag{B.7}$$

Applying (B.5) and (B.7) to (B.2),

$$r(t) = \{\cos(2\pi t f_L)\cos[2\pi t\Delta f(N-1)/2] - \sin(2\pi t f_L)\sin[2\pi t\Delta f(N-1)/2]\}\frac{\sin[2\pi t\Delta f(N/2)]}{\sin(2\pi t\Delta f/2)}$$

$$= \cos[2\pi t(f_L + \Delta f(N-1)/2)]\frac{\sin(\pi t\Delta fN)}{\sin(\pi t\Delta f)}$$

$$= \cos(2\pi t f_c)\frac{\sin(\pi t\Delta fN)}{\sin(\pi t\Delta f)}. \qquad \text{(B.8)}$$

where $f_c = f_L + \frac{1}{2}\Delta f(N-1)$, the average frequency of the group of sinusoids. Eq. (B.8) is equivalent to (3.24), completing the derivation.

Two remarks need to be made regarding the ratio on the right of (B.8). First, this expression is sometimes called the "discrete-time sinc function," since it often appears in discrete-time problems. It is analogous to the sinc function (i.e., sin[x]/x) that arises in continuous-time situations. Like the traditional sinc function, the expression has singularities that are easily resolved with L'Hopital's rule. The singularities occur when $t\Delta f$ is an integer. In these instances the ratio converges to the following:

$$\lim_{\substack{t\Delta f \to k \\ k \in \{\text{integers}\}}} \frac{\sin(\pi t\Delta fN)}{\sin(\pi t\Delta f)} = \begin{cases} -N, & N \text{ even \& } k \text{ odd}, \\ N, & \text{otherwise}. \end{cases} \qquad \text{(B.9)}$$

Second, in this thesis, N is always odd, and therefore this function is periodic with period $1/\Delta f$. In situations where N is even, this function is periodic with period $2/\Delta f$, but one can say it is *nearly* periodic with period $1/\Delta f$, since the difference is just an algebraic sign. The magnitude has period $1/\Delta f$ regardless of the parity of N.

REFERENCES

[Bal81] Balabanian, N., T. Bickart, *Linear Network Theory: Analysis, Properties, Design and Synthesis*, Matrix (Beaverton, OR: 1981).

[Bal86] Balaberda, R., S. R. Mishra, "On the Use of the HP-8510 Network Analyzer for Antenna Pattern Measurements," Proceedings of the AMTA, October 1986, pp. 13-17.

[Bee86] Beeckman, P. A., "Millimeter-Wave Antenna Measurements With the HP 8510 Network Analyzer," *Microwave Journal*, vol. 29, no. 7, July 1986, pp. 89-102.

[Boy85] Boyles, J. W., "Measuring a Dipole Antenna Radiation Pattern Using Time Domain and Gating," *RF Design*, September 1985.

[Bur73] Burrell, G. A., and A. R. Jamieson, "Antenna Radiation Pattern Measurement Using Time-to-Frequency Transformation (TFT) Techniques," *IEEE Trans. on Antennas and Propagation*, vol. AP-21, September 1973, pp. 702-704.

[Che89] Cheng, David K., *Field and Wave Electromagnetics*, Addison-Wesley (Reading, Mass.: 1989).

[Cou90] Couch, L. W., *Digital and Analog Communication Systems*, 3rd edition, Macmillan (New York: 1990).

[Dom87] Dominek, A. K., L. Peters, Jr., and W. D. Burnside, "A Time Domain Technique for Mechanism Extraction," *IEEE Trans. on Antennas and Propagation*, vol. AP-35, March 1987.

[Eva85] Evans, Gary, "Far Field Correction for Short Antenna Ranges," Proceedings of the AMTA, October, 1985, pp. 34-1 to 34-9.

[Eva90a] Evans, Gary, *Antenna Measurement Techniques,* Artech House (Norwood, Mass: 1990).

[Eva90b] Evans, Gary, *Antenna Measurement Techniques,* Artech House (Norwood, Mass: 1990), Chapter 8.

[Hen89] Henderson, A., J. R. James, P. Newham, G. Morris, "Analysis of gating errors in time domain antenna measurements," *IEE Proceedings*, part H, vol. 136, August 1989, pp. 311-320.

[Hes88] Hess, D. W., and V. Farr, "Time Gating of Antenna Measurements II," *Proceedings of the AMTA*, Atlanta, GA, September 1988, pp. 1-9 to 1-18.

[Kun90] Kunath, R. R., M. J. Garrett, "Near-Field Antenna Testing Using the Hewlett Packard 8510 Automated Network Analyzer," *Proceedings of the AMTA*, Philadelphia, October 1990.

[Lam87] Lambert, K. M., R. C. Rudduck, and T. H. Lee, "Far-Field Pattern Measurements and Time Domain Analysis of Reflector Antennas in the Compact Range," *Proceedings of the AMTA*, Seattle, WA, October 1987, pp. 44-49.

[Lam71] Lamensdorf, D., and L. Susman, "An analysis of Some Directive Antennas Using Time Domain Measurements," *1971 G-AP International Symposium*, Los Angeles, CA, September 1971.

[Lee88] Lee, T. H., and W. D. Burnside, "Cross Range Processing of Patterns for Large Reflector Antennas to Obtain Radiation Centers," *Proceedings of the AMTA*, Atlanta, GA, September 1988, pp. 1-19 to 1-24.

[Mar90] Maricevic, Z., T. K. Sarkar, Yingbo Hua, "Time domain antenna measurements with HP 8510 using matrix pencil method," *Symposium on Antenna Technology and Applied Electromagnetics 1990 Conf. Proceedings*, Winnipeg, August 1990, pp. 301-317.

[Mar87] Markman, D. J., R. E. Hartman, "An automated antenna measurement system using the HP 8510," *IEEE Instrumentation and Measurement Technology Conf.*, Boston, April 1987, pp. 182-185.

[Men84] Mensa, D. L., "Extraction of Narrowband Responses from Wideband RCS Data," *Proceedings of the AMTA*, October 1984, pp. 2B5-1 to 2B5-23.

[Nat94] *LabVIEW User Manual for Windows*, National Instruments (Austin, TX: 1994).

[Nic68] Nicolson, A. M., "Broadband Microwave Transmission Characteristics from a Single Measurement of the Transient Response," *IEEE Trans. on Instrumentation and Measurement*, vol. IM-17, December 1968, pp. 395-402.

[Opp89] Oppenheim, A. V., and R. W. Schafer, *Discrete-time Signal Processing*, Prentice-Hall (Englewood Cliffs, NJ: 1989).

[Poz84] Pozar, D. M., and D. H. Schaubert, "Scan Blindness in Infinite Phased Arrays of Printed Dipoles," *IEEE Trans. on Antennas and Propagation*, vol. AP-32, June 1984, pp. 602-610.

[Pra86] Pratt, T., and C. W. Bostian, *Satellite Communications*, John Wiley (New York: 1986).

[Pre95] Predoehl, A. *Range Runner User's Guide*, Antenna Group of Virginia Tech (Blacksburg, VA: 1995).

[Seb89] Sebesta, R. W., *Concepts of Programming Languages*, Benjamin/Cummings (Redwood City, CA: 1989).

[Ser84] *Series 1780 Programmable Microwave Receiver: Operator's Manual*, third edition, Scientific Atlanta (Atlanta, GA: 1984).

[Stu81] Stutzman, W. L., and G. A. Thiele, *Antenna Theory and Design*, John Wiley (New York: 1981).

[Wai93] Waiyapattanakorn, C., and C. G. Parini, "Theoretical and Experimental Investigations of Using Time Domain Gating in Antenna Pattern Measurements," *Eighth International Conf. on Antennas and Propagation* (IEE), vol. 1, London, April 1993.

[Zeg95] Zeger, A. E., and B. S. Adams, "Pattern Measurement of Ultralow Sidelobe Level Antennas," *Proceedings of the AMTA*, November 1995, pp. 43-48

VITA

Andrew M. Predoehl was born in Pendleton, Oregon, and grew up in Olney, Maryland. He graduated from high school in 1987 and began studying computer engineering at Virginia Tech in the fall of that year. In 1992 he completed his bachelor's degree and began graduate school, joining the Satellite Communications Group as a graduate research assistant. There he worked on the digital receiver / detector for the Advanced Communications Technology (ACTS) Propagation Terminals until 1993. This was followed by research on frequency-locked loop and phase-locked loop digital receivers. After this he worked as a graduate research assistant on designing and developing software for the Virginia Tech rooftop antenna range, followed by development of a proof-of-concept switched-beam adaptive antenna system for cellular base stations. He intends to work as a design engineer in the field of digital signal processing.