Early Advances in Radar Technology for Aircraft Detection

Donald L. Clark

In its early years, Lincoln Laboratory developed critical components of an air-defense system to guard North America against the threat of intercontinental bombers carrying nuclear weapons. Lincoln Laboratory used digital computer technology to automate several functions of the air-defense system and improve the quality of digitized radar data processed by the air-defense system. This article describes some of the experimental and theoretical efforts that led to early advances in radar technology for aircraft detection.

PROJECT LINCOLN, THE FORERUNNER of Lincoln Laboratory in Lexington, Massachusetts, was initiated in 1951 to address the problem of defending the continental United States and Alaska against intercontinental bombers. Researchers faced the challenge of applying advanced technology to achieve the following improvements in the air-defense system: (1) consolidated command and control at a central post in each air-defense sector of about 100,000 square miles, (2) provided coverage against low-flying aircraft by supplementing the principal long-range radars in each sector with numerous short-range gap-filler radars, (3) automatically transferred filtered data from each radar to its central command center, and (4) improved communication between each command center and its interceptors. (Reference 1 provides a more extensive account.)

This was an exciting era in the life of the Laboratory, with a talented and highly motivated staff, a can-do spirit, and minimum administrative formality. An Air Force unit at nearby Hanscom Field provided substantial logistic support. Project Lincoln set up an experimental air-defense sector called the Cape Cod System in southeastern New England, as shown in Figure 1. The initial long-range radar for the Cape Cod System, the AN/FPS-3, was an L-band radar with a nominal range of 200 miles on a high-flying bomber. Low-flying aircraft could evade the coverage of the AN/FPS-3 by staying below its horizon. Gap-filler radars, assembled mostly from World War II components, operated at S-band with a nominal range of 32 miles. Later, the Cape Cod System was extended to include additional long-range radars at Montauk Point in Long Island and at West Bath, Maine. These radars supplied only range and azimuth coordinates. The heights of designated targets were measured separately by a small number of height-finder radars.

Data from all the radars were transmitted over ordinary leased telephone lines to a command center in Cambridge, Massachusetts, where the Whirlwind computer, and later the system prototype AN/FSQ-7 in Lexington, Massachusetts, processed the data in real time to track aircraft, assist operators to perform command functions, and guide manned interceptors. The defense system that grew out of this effort was called the Semi-Automatic Ground Environment (SAGE). Developing SAGE was the major activity during Lincoln Laboratory's early years [2].

The SAGE system also engaged Lincoln Labora-

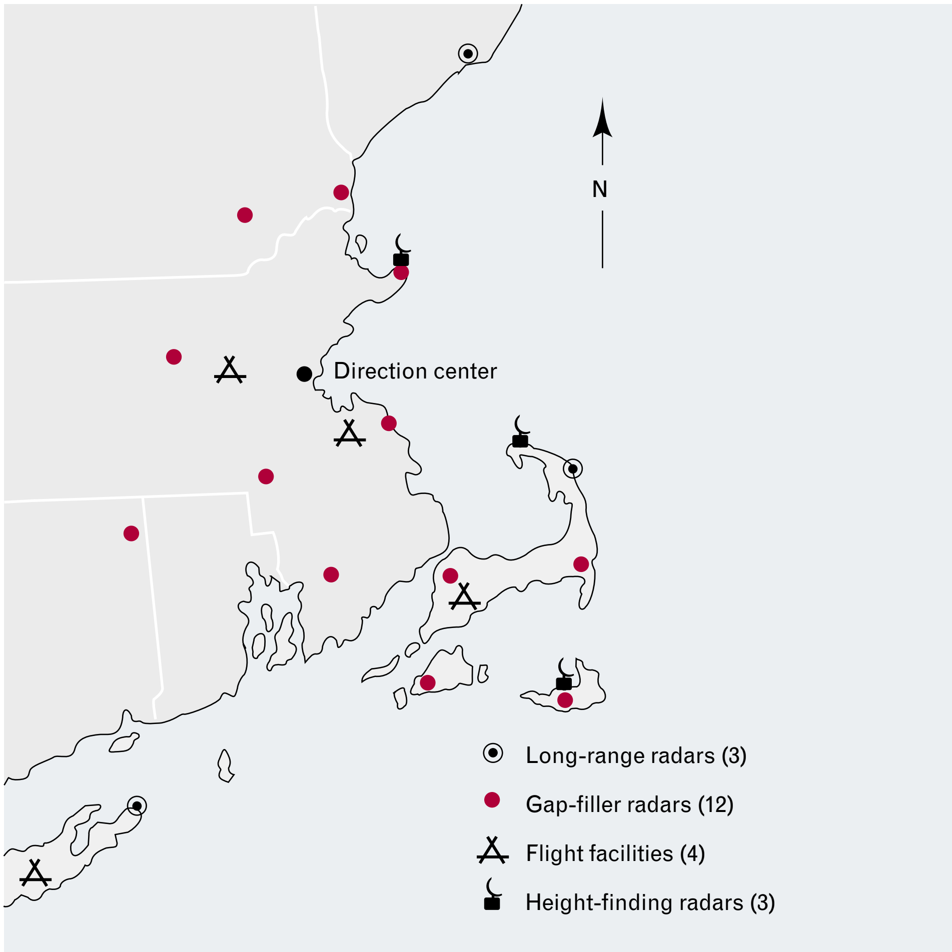

FIGURE 1. Locations of radar sites in the Cape Cod System.

tory in related air-defense problems. For example, in parallel with developing the Cape Cod System, Lincoln Laboratory helped design and develop the Distant Early Warning (DEW) Line of radars across northern Alaska, Canada, and Greenland. The need for highly automated detection during operations in the harsh environment of the far north complicated the design criteria for the DEW Line radars. The early operations of the DEW Line and the Cape Cod System raised fundamental questions about radar signal processing within an integrated, digitized system. This article reviews the theory and practice of the principal problems in the collection, filtering, transmission, and centralized processing of digitized radar data. A significant component of the DEW Line is described in the sidebar by Edwin L. Key entitled "The Sentinel Radar."

Automated Data Transmission

To detect low-flying aircraft, a radar cannot avoid illuminating the surface of the earth, which produces clutter, or echoes, from many reflectors other than aircraft. Distinguishing between clutter and the signals from aircraft relies on filtering out as much clutter as possible. Effective filtering uses the Doppler effect and therefore requires a high degree of frequency stability in the transmitted signal. A radar must be resistant to incidental interference and intentional jamming. Furthermore, the use of automated detection, data transmission, and subsequent computer processing imposes requirements on signal detection and filtering different from those for use with a display for a human operator.

Before Project Lincoln was initiated, the Air Force

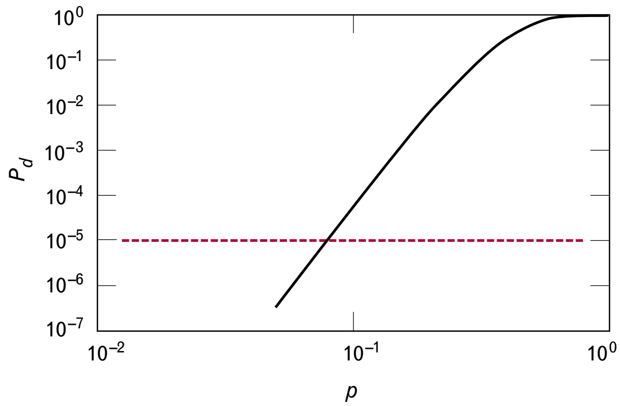

FIGURE 2. Probability of detection P_d at the output of a digital integrator versus the probability p of a threshold crossing on a single trial. The sliding window is sixteen units long. Eight or more threshold crossings constitute a detection. A useful detection probability of 0.9 can be attained with those values when p equals 0.625 or greater.

Cambridge Research Laboratory in Cambridge, Massachusetts, developed a method called Slowed-Down Video (SDV) that serves in this discussion as a generic model of automated detection. In SDV, range-gated data from a scanning radar are digitized with a single binary digit per range gate. A one value represents a threshold crossing by the signal plus noise and a zero value represents no crossing. Within each range gate the ones were counted in a sliding window as the radar beam scanned a target. At typical scan rates the radar beam stayed on a target long enough for a dozen or more pulse returns from that target. A sufficient number of ones within a window represented a target detection and triggered a detection signal that was transmitted over the telephone line. Digits representing detections or nondetections for every range gate and every beamwidth were transmitted sequentially within the bandwidth of the telephone line. This process yielded data at the receiving end suitable for computer processing and generating a digitized version of the radar display. (See also the article entitled "Radar Signal Processing," by Robert J. Purdy et al., in this issue.)

A form of digital signal integration can be simply characterized for this discussion by considering a sliding window observing a single range gate. A counter totals the threshold crossings within the window. A detection occurs when a designated number of threshold crossings is counted. We need to know the probability of a detection in the window as a function of the probability of a threshold crossing in each range gate. In mathematical terms, the sliding window is m units long and k or more threshold crossings constitute a detection. J.V. Harrington, who analyzed this process, used the fact that the probability of detection P_d is the sum from k to m of the well-known binomial distribution $b(m, k, p)$, where p is the probability of a threshold crossing in a single trial [3].

Figure 2 shows P_d as a steep function of p for representative values of m and k. A useful detection probability of 0.9 can be attained with those values when p equals 0.625 or greater. Unfortunately, threshold crossings on noise alone can also trigger false detections. If the threshold-crossing-on-noise probability p rises much above 0.08, then the number of false alarms can greatly outnumber true detections, overloading the telephone line and computer with useless data. A technique used at the time to protect the computer was simply to map out and excise any area with too many false targets, thereby losing radar coverage in that area. Clearly, there was a need to provide the cleanest data possible, which translates into keeping the false-alarm rate low. Although the tolerable false-alarm probability P_{fa} is not a sharply defined quantity, in the following discussion we use a P_{fa} value of 10^{-5}, which would average less than one false target per scan for a typical system that has a few tens of thousands of range-azimuth cells per antenna scan. This level corresponds to a threshold-crossing-on-noise probability p of about 0.08 and is acceptable because the system should be able to accommodate several tens of apparent aircraft returns per scan.

Theoretical Work

Wilbur B. Davenport, Edward J. Kelly, Irving S. Reed, William L. Root, Irwin Shapiro, and Richard P. Wishner led the extensive theoretical work done to establish and understand fundamental limits on radar detection, filtering, and parameter-estimation capabilities [4, 5]. They worked individually or collaborated in various ways. Reed, a mathematician who professed to have little understanding of practical problems, circulated widely around the Laboratory, talking with engineers about their radar problems.

THE SENTINEL RADAR

Edwin L. Key

In the early 1950s, the United States decided to deploy a line of radars across northern Alaska, Canada, and Greenland to provide early warning of a possible USSR bomber attack on North America. This warning system was the Distant Early Warning (DEW) Line. From the outset, the logistical support of radars in such remote locations was problematic and expensive. Because air traffic in these extreme northern regions was light, however, the architects of the DEW Line felt that constant observation of radar screens by human operators was unnecessary. Consequently, they planned for the radars to be unattended except when aircraft were actually penetrating the warning zone. This arrangement was achieved by providing automatic alarms to alert operators when aircraft detection occurred. Upon such alerts, the operators could monitor the radar displays and evaluate the observed circumstances for potential threats. Since the operators did not need to constantly attend the radar, they were largely free to perform routine site duties that would otherwise require additional personnel.

Table A. Parameters of the Sentinel Radar

RF frequency	570–630 MHz
Peak power	150 kW
Average power	3 kW
Pulse length	40 μsec (detection) 5 μsec (threat analysis)
Pulse compression	
Barker 13-segment code	39 μsec
compressed to	3 μsec
Transmitter output tube	Klystron with 62-dB gain
Pulse-repetition frequency	500 pulse/sec
Receiver noise figure	<6.5 dB
Antenna aperture	45 ft × 25 ft

Lincoln Laboratory designed the experimental automatic alarm system that was used on the modified AN/TPS-1D radar, which later became the AN/FPS-19. (See the article entitled "Distant Early Warning Radars: The Quest for Automatic Signal Detection," by F. Robert Naka and William W. Ward, in this issue.) Because the design was for an existing radar with parameters not optimum for the purpose, the results were not entirely satisfactory. These shortcomings motivated Herbert G. Weiss to design and advocate a new radar that better served the requirements for automatic detection. The concept was approved, and in 1954 Lincoln Laboratory began the development. The new radar, called Sentinel, was completed and went on the air for testing at Lexington in 1955. It incorporated in a unified system design many of the innovations that are described in the accompanying article. Table A shows the major parameters of the Sentinel radar.

The Sentinel radar incorporated several unusual features for that time. These novel features resulted from the requirement that it be essentially unattended. The radio-frequency (RF) power source was a high-gain four-cavity klystron amplifier, which employed a special modulating anode to pulse the beam. The klystron in conjunction with two very stable oscillators provided high-quality coherence for clutter cancellation and velocity filtering.

The crystal-controlled RF (540 to 600 MHz) stable local oscillator (STALO) was the rock on which Sentinel's frequency stability rested. The coherent local oscillator (COHO) was a crystal-controlled 30-MHz oscillator. Its output was added to the STALO's output to provide transmitter excitation (570 to 630 MHz). The COHO signal was also used as the phase reference for filtering the echo signals after they had been downshifted to an intermediate frequency of 30 MHz plus Doppler shift by mixing the RF echoes with the STALO signal.

The pulse-repetition frequency (PRF) was 500/sec to provide a large unambiguous velocity range for velocity filtering to reject returns from migrating birds that were a problem in the Arctic. The PRF limited the unambiguous range to 162 nm, but since the purpose of the radar's operation was to provide a "trip-wire"-like warning, this was not a concern. Within this limited unambiguous range Sentinel had a substantial detection margin to allow for degradation. The pulse length for normal warning operation was 40 μsec to provide high pulse energy for detection and to completely cover the unambiguous range with 50 range gates. After detection the pulse length could be reduced to 5 μsec for threat analysis. The Sentinel radar was an early application of pulse compression, which had both theoretical and practical significance.

It was recognized that the 40-μsec pulse would result in rather large clutter power within a range-azimuth resolution cell, but the corresponding transmitted pulse energy was required for detection performance. Introducing phase-coded compression similar to that used in the AN/FPS-17 radar allowed the radar to transmit a long pulse while achieving range resolution that corresponded to a short transmitted pulse. (For a discussion of the AN/FPS-17 radar, see the article entitled "Radars for the Detection and Tracking of Ballistic Missiles, Satellites, and Planets," by Melvin L. Stone and Gerald P. Banner, in this issue.) A simple Barker 13-segment phase-reversal code was designed for the Sentinel radar to test whether pulse compression could reduce clutter. The pulse-compression waveform consisted of a 39-μsec pulse with 3-μsec subpulses coded with a particular sequence of phase reversals, providing 13-to-1 compression. Once the clutter-reduction claims were verified, researchers included pulse compression in the Sentinel radar's design.

The Sentinel radar was acquired by the Air Force, renamed the AN/FPS-30, and manufactured by Bendix. The AN/FPS-30 was deployed by the Air Force in the extension of the DEW Line across southern Greenland. It was reported that the pulse compression proved to be the savior of the system. Sometimes a large ice floe off the coast of Greenland produced clutter returns with a 40-μsec pulse that exceeded the subclutter-visibility capabilities of the radar. However, the 13-to-1 compression provided enough clutter reduction for the radar to be able to see targets.

Edwin L. Key joined Lincoln Laboratory in 1951 and contributed to a variety of radar programs, including SAGE, the DEW Line, BMEWS, and the Millstone Hill facility. He transferred to the MITRE corporation upon its formation in 1959. At MITRE, he held positions of increasing responsibility that culminated in his position as senior vice president for research and engineering. Although retired, he continues to consult in the areas of radar systems, radar technology, signal processing, and communication systems.

These discussions often gave him the idea for an analysis that he could perform. Then, in useful cross-fertilization, he would report back on his analytical results to the engineer who had inspired the idea. He was one of the more prolific authors of technical reports in the early days of the Laboratory. One of his major contributions to the digital-processing community is the origination of the Reed-Solomon error-correction algorithm, which is discussed at the end of this article.

A significant problem under investigation during this time involved signal integration. Our earliest knowledge of signal integration came from Ruby Payne-Scott [6], J.I. Marcum [7] and John V. Harrington [3]. Confusion existed about the relative advantages of noncoherent integration and coherent integration. Over a period of time, with contributions from several people, we learned that coherent integration can offer a substantial advantage, especially for signal-to-noise ratios near unity or lower. The advantage is greatest when we know precisely the frequency of the signal being integrated. The advantage is reduced somewhat when the signal can have a range of frequencies, as with Doppler frequencies from moving targets with unknown velocities.

Another problem at the time was establishing how accurately a radar could estimate various target parameters such as range, velocity, acceleration, and azimuth and elevation angle. Several people within the Laboratory, including Kelly, Roger Manasse, Reed, and Root, and others outside the Laboratory, including P. Swerling, addressed the problem. Their work, and that of others, is well summarized by Swerling in chapter 4 of Merrill I. Skolnik's *Radar Handbook* [8].

Chaff, consisting of many small scatterers, was a countermeasure used widely in World War II to confuse enemy radars. Kelly treated extended targets like chaff and precipitation with a mathematical model for the radar echo as coming from a random collection of scatterers, necessarily highly idealized [9].

The characteristics of the chaff radar echo depend upon the distribution of the scatterers, their bulk motion, and their motion relative to each other, as well as on the characteristics of the radar pulse that illuminates them. The radar echo from many scatterers has noiselike qualities and coherent qualities. Depending on the wind, the Doppler frequencies can be high enough to seriously compromise Doppler-filtering schemes. Wind shear or turbulence can generate a spread of Doppler frequencies that further complicates the job of filtering. Changing the radar frequency randomly from pulse to pulse can destroy the coherence of the echo. (A way of exploiting this fact is briefly described later.) Other sources provide a more detailed account of this theoretical work [9–13].

Looking to the future, well beyond air defense, Shapiro studied the ability of radars to predict ballistic missile trajectories. His monograph proved to be the defining work on this subject [13]. Shapiro then turned his attention to radar astronomy, and he has had a distinguished scientific career at MIT and at the Smithsonian Astrophysical Observatory.

Characterization of Clutter

The author's acquaintance with radar clutter began in the summer of 1952 under the guidance of Robert C. Butman, an experienced radar engineer. We had the use of an experimental S-band radar located on the roof of MIT's tallest building. It had a good view of the area surrounding Cambridge, which included Logan Airport and Hanscom Field. We spent many hours observing air traffic and clutter, and experimenting with the Doppler-filtering moving-target indicator (MTI). In principle, MTI filtered out clutter signals from stationary and near-stationary reflectors while passing signals from aircraft having significant radial velocities with respect to the radar. A rough but useful judgment of the ability of the MTI to reject clutter could be made by comparing a plan position indicator (PPI) display of the unfiltered radar scene with a display produced by the MTI. Much less obvious was the radar's ability to follow an aircraft through a heavily cluttered area. At one point Butman, taking advantage of Air Force logistic support, arranged for a flight test to start learning something about that ability. At the appointed time a trainer aircraft appeared and circled our site. Butman established contact by using a war-surplus radio. He requested the pilot to fly over a heavily cluttered area northwest of the radar to see how well the MTI could pick the aircraft out of the clutter. When the aircraft was lined up on the desired course he asked the pilot to continue flying in that direction until he received further instructions. At that point our radio emitted clouds of smoke and died. Such was the author's introduction to flight tests.

The twelve gap-filler sites of the Cape Cod System provided an opportunity to study a variety of clutter. An instrument called the MTI Site Evaluator was constructed and brought to each site to observe clutter in conjunction with the MTI. Unfortunately, the name of the device initially terrified some of the site technicians. The instrument mapped areas where aircraft might be obscured by strong background clutter. After the survey of all sites, the site with the largest and most intensely obscured area was selected for a flight test. A medium bomber, following radio instruction from the site, was tracked on the radar dis-

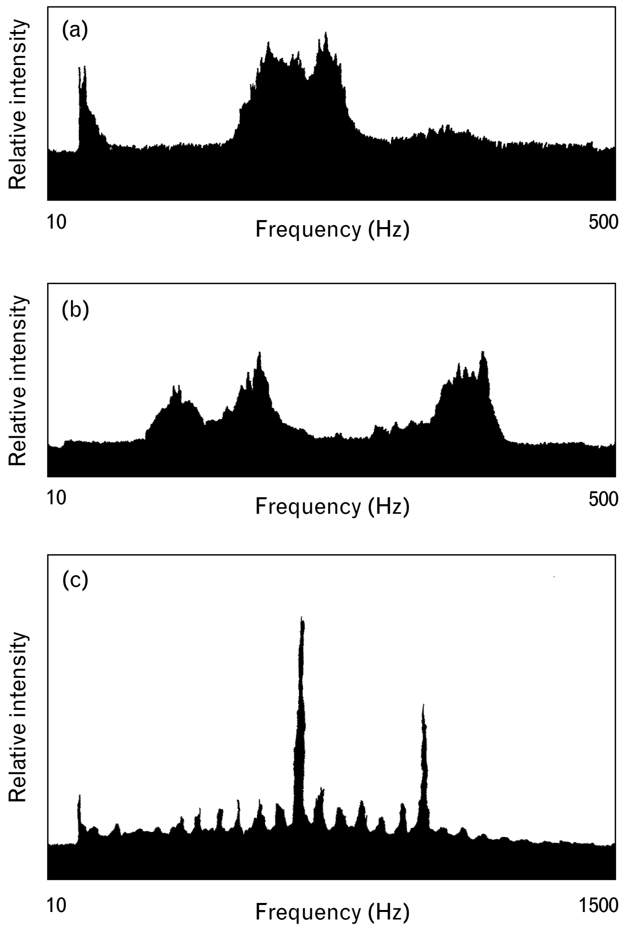

FIGURE 3. (a) Doppler spectrum of rainstorm echoes observed with an S-band radar. (b) Doppler spectrum of chaff echoes observed with an S-band radar. (c) Doppler spectrum of a C-45 aircraft inbound observed with an S-band radar.

play as it flew over what appeared to be the most severely obscured area. It proved difficult to find an area where the echo from the bomber was obscured for more than a scan or two. This finding was consistent with observations of aircraft targets at other sites. At the time, we concluded that very intense clutter echoes were due to specular reflections from fixed targets of limited areal extent, such as tall buildings and water towers. If filtered out by the MTI they were not likely to negate the ability of the radar to track large aircraft. The term interclutter visibility was later used to describe this ability to track objects in areas of strong background clutter.

Some limited observations were made with an S-band radar on echoes from rain and chaff. Observations consisted of numerous Doppler spectra measured within a single range gate with the antenna pointed to the region of interest. Figures 3(a) and 3(b) show sample spectra from rain and from chaff, respectively. Both spectra were found to be highly variable, depending upon wind conditions. A spectrum from a C-45 aircraft inbound is shown in Figure 3(c). Either rain or chaff could have Doppler frequencies overlapping those expected from aircraft, thereby seriously complicating the task of Doppler filtering.

These limited observations of clutter left much to be desired. They provided only a qualitative basis for designing filtering schemes. A proper characterization of clutter awaited another era, when low-altitude, low-cross-section cruise missiles were the driving concern. (See the article entitled "Radars for the Detection and Tracking of Cruise Missiles," by Lee O. Upton and Lewis A. Thurman, in this issue.)

Clutter Filtering

The principal technique initially available for dealing with radar clutter was MTI. In this scheme, the radar's phase-detected video signal was delayed by one interpulse interval and subtracted from the next video signal. The signal from a stationary target such as a water tower would be effectively canceled. The signal from a target having a significant radial velocity with respect to the radar would be Doppler-shifted and, in general, would not cancel. Chapter 17 of Skolnik's *Radar Handbook* offers a useful discussion of this type of MTI [14]. During our experiments, the signal was delayed as a sound wave propagating through a column of mercury driven by a piezoelectric transducer. With careful adjustment, sharp nulls could be attained on stationary targets, but they were hard to maintain due to temperature changes and other instabilities of the delay line. Effort was directed at developing better means of delaying the video signal. After considerable trial and error, the best approach was found to be the use of a delay line in which the sound wave followed a folded path within a slab of fused quartz. A cancellation null depth of 37 dB was eventually demonstrated, with excellent stability.

The filter response of this type of MTI had the form of a rectified sine wave with peaks at odd multiples of half the pulse-repetition frequency (PRF) and nulls (blind speeds) at multiples of the PRF. (A blind speed occurred when the target moved radially an integral number of half wavelengths in one interpulse period.) Other approaches to Doppler filtering were tried that offered some flexibility in shaping the filter response. In one approach, coherent video was range-gated and filtered with an analog filter for each range gate. In another approach, Thomas C. Bazemore and Bruce Nelson saved the polarities of the range-gated coherent video as a sequence of binary digits in a shift register. A diode or resistor array attached to the shift register allowed periodicities that correspond to Doppler frequencies to be detected. In effect one had a bank of elementary digital filters in a sliding window. This scheme was tested with a few range gates on an S-band radar. Although it showed promise, the digital technology available at that time required one vacuum tube flip-flop per binary digit, which made a full-scale implementation impractical.

John P. Perry and the author investigated another approach to Doppler filtering called "Sinufly," a technique borrowed from C.W. Sherwin at the Coordinated Science Laboratory of the University of Illinois. This technique used a storage tube similar to a cathode-ray tube except that an electrostatic storage surface replaced the phosphor on the face of the tube. The electron beam, modulated by radar video, was scanned in a raster across the storage surface, laying down a pattern of electric charge. When the surface was completely scanned, the unmodulated beam was scanned in a raster at right angles to the first to read out the stored electric charge. This arrangement had the effect of range-gating the video and allowing the Doppler frequencies, multiplied by a large factor, to be read out sequentially, range gate by range gate. By switching between two tubes all of the radar video could be captured. With this scheme a single bank of filters was shared among all range gates, which permitted experimentation with various filtering schemes. The system had many attractive features but was limited by the inadequate dynamic range of the storage tubes that we used. The tubes were laboratory samples that were never further developed; hence we had to abandon this scheme.

Although we explored a variety of ideas for clutter filtering, we found that many sophisticated approaches could not be reasonably supported by the memory technology available at the time [15]. The best memory available to us in this era was the quartz delay line. A demonstration system exploiting quartz delay lines is described briefly below. A small magnetic-core memory, which provided some interesting possibilities, arrived too late to be fully exploited.

Jamming and Interference

Andrew Bark and Robert Bergemann, working on radar countermeasures, demonstrated with a jammer developed during World War II the radar problems caused by jamming or by incidental interference. The jammer consisted of a mechanically tunable continuous-wave (CW) magnetron that radiated through a small horn antenna in the tail of a test aircraft. Its frequency could be swept rapidly back and forth across a broad band and its signal was sufficiently strong to penetrate the sidelobes of the radar antenna as well as the main lobe. When swept through the radar frequency band, it produced a strong signal that cluttered the display and reduced the radar's sensitivity. With automated detection, it overwhelmed the system with false alarms.

A remedy, suggested by Robert H. Dicke during an earlier study, was to limit or clip the amplitude of a radar signal. This technique became known as the Dicke Fix. (Actually, the idea of clipping signal amplitude to reduce the effects of interference was a new application of an old idea. Radio hams had already been reducing interference effects in their receivers this way.)

We explored several methods for clipping the amplitude of a signal. The simplest, and one of the most effective variations, was to use a wideband intermediate-frequency (IF) amplifier with a hard limiter followed by a narrowband filter and rectifier. A hard limiter amplified the receiver noise to drive the limiter to saturation, so that there was negligible variation of the amplitude at its output with or without signal or interference present. The narrowband filter responded more strongly to the radar signal than to noise, allowing aircraft to be detected. In this context, narrowband meant a bandwidth matched to the radar pulse; wideband meant bandwidth several (e.g., ten) times greater.

A receiver using a technique like this was found by experiment to be much more resistant to jamming and interference than a more conventional non-limiting receiver. A major frustration at the time was the lack of an analytical model to describe the effects we were seeing experimentally. Such a model was developed some forty years later and is described in the next section.

Analytic Performance of a Hard-Limiting Receiver

An example using a simple analytic model of a hard-limiting receiver can help us to understand the hard-limiting receiver and to compare it with a conventional linear receiver. In this model, a wideband filter precedes the narrowband filter, where the ratio of bandwidths is n. The narrowband filter is thus presented with n independent samples of noise passed by the wideband filter, plus the signal (if present). The noise is modeled as having random phase over 360° and a Gaussian amplitude distribution with zero mean and unit variance. The narrowband filter, tuned to the signal frequency, adds the n samples. The noise adds noncoherently; the signal, having constant phase (modulo 2π), adds coherently. In the linear receiver a threshold is set such that when the threshold is exceeded, the receiver puts out a one, otherwise, a zero. In the hard-limiting receiver, each of the n samples is limited to unit amplitude, with phase the same as that of the signal plus noise of the linear receiver. We can picture these samples as n unit vectors that add in random-walk fashion when noise predominates, and that line up to add in phase when signal predominates. Again, a threshold is set to produce a one when exceeded and a zero otherwise. (For an analytical discussion of related ideas see Reference 16.)

This model was implemented in a program that generated n samples of signal plus noise, as described. The resulting n vectors were added for each receiver, and the amplitude of the result was compared to a fixed threshold for each. By repeating this process many times Monte Carlo fashion, we could estimate the probability p of detection/false alarm per range gate. Using the binomial formula referred to earlier we could calculate the probability of (apparent) detection P after digital integration.

Interference and/or jamming were modeled as a

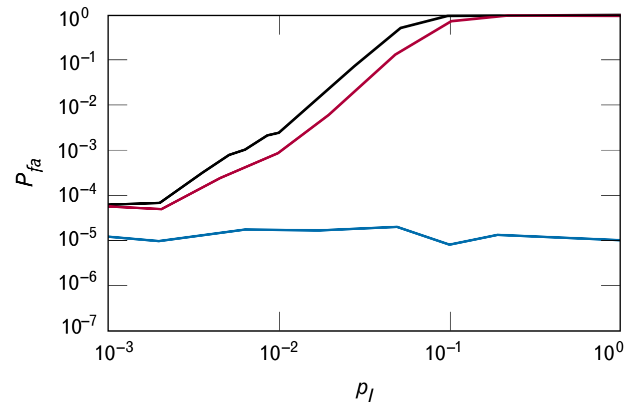

FIGURE 4. Probability of false alarm P_{fa} versus probability of interference p_I. The black curve represents data from a linear receiver with interference 100 times the noise. The red curve represents data from a linear receiver with interference 10 times the noise. The blue curve represents data from a hard-limiting receiver with interference 10 times the noise. Only the hard-limiting receiver maintains a P_{fa} value of 10^{-5} for all values of p_I.

probability p_I that a noise sample would be taken from a distribution with variance I many times that of the noise, but otherwise similar. This probability p_I could be chosen anywhere in the range from 0 to 1.

Figure 4 shows the behavior of the two types of receiver with fixed thresholds set to produce a probability p of about 0.08 for a threshold crossing on noise alone. As p_I was increased, the false-alarm probability P_{fa} was observed for values of I that are 10 and 100 times the noise. Figure 4 shows that the hard-limiting receiver maintained P_{fa} at about 10^{-5} for all values of p_I. It is clear that the linear receiver produced intolerable false-alarm probabilities by the criterion described earlier in this section for p_I greater than about 10^{-2}. The curves shown are somewhat irregular, due to the limited statistics of the Monte Carlo model.

The threshold for the linear receiver can be adjusted to keep its false-alarm probability nearly constant as p_I is changed. The threshold has to be raised substantially as p_I increases to maintain the false-alarm probability, thereby desensitizing the receiver. We can then observe the probability of detection with a signal present. Figure 5 shows P_d as a function of p_I for a signal whose peak amplitude was the square root of 10 times the variance of the wideband noise. The curves shown are for I equal to 100 times the noise. It is apparent that the linear receiver lost its ability to

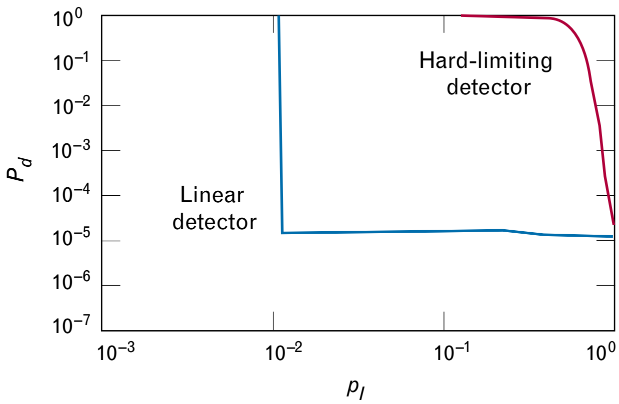

FIGURE 5. Probability of detection P_d at the output of a digital integrator versus the probability of interference p_I. Interference was 100 times noise. Peak signal amplitude was the square root of 10 times the wideband noise variance.

detect the signal at small values of p_I. The hard-limiting receiver maintained a useful detection probability for values of p_I up to about 0.4, for the combination of parameters chosen.

To summarize, in the presence of intense interference the linear receiver either produced too many false alarms or became desensitized. The hard-limiting receiver, by contrast, maintained a constant low false-alarm probability, together with useful sensitivity in the presence of severe interference. Franklin A. Rodgers, who briefly led the group in which this work was done, described the class of receivers in our experiments as constant false-alarm rate (CFAR), which was also a play on Rodgers's initials.

Performance of Actual Receivers

Figure 6 shows a practical comparison of two actual receivers subjected to jamming. Figure 6(a) shows a time exposure of an output display of a linear receiver subjected to severe jamming from three airborne jammers on a single plane. Figure 6(b) shows a time exposure of an output display of a hard-limiting receiver under the same conditions. The hard-limiting receiver in this case used wideband video with a zero-crossing counter. The lines of blips represent aircraft tracks. The results are consistent with the results of the analytical model described above.

Some time after we had achieved the results briefly described above, we learned that D. Griffin at Harvard University had made some interesting observations of the ability of bats to rely on their sonar to navigate and capture insects. In particular, he had tested bats in the presence of interfering noise, and they appeared to have an astonishing capability to resist the noise. Had we overlooked something that the bats could teach us?

The Laboratory arranged to have J.J. Gerald McCue, assisted by David A. Cahlander, work with Griffin to follow up on his observations in more detail. They set up a carefully instrumented enclosure in which the bats could fly freely. Instrumentation included a strobe light and high-speed movie camera to photograph the bats in flight, a microphone and recording system to record their chirps, and an adjustable noise generator that filled their enclosure with continuous near-white noise to jam their sonar. Movies of flying bats were played back in slow motion. Recorded chirps were synchronized with the movie and slowed down sufficiently to bring the chirps within human audible range. The movies and sound provided a graphic and convincing way of observing the bats. It was an enjoyable project whose propensity to grow had to be restrained. We learned that, alas, the bats had no more ability to resist jamming than could be accounted for by existing models [17–19].

Master-Oscillator/Power-Amplifier Transmitter

The earliest microwave radars used magnetrons to generate the power required for the transmitter. Magnetrons have an honorable history as the invention that made possible the development of microwave radar during World War II. Magnetrons operated as self-excited oscillators whose characteristics, however, left something to be desired. Early in the life of Lincoln Laboratory an arrangement was worked out with the Physics Department at Stanford University and with Varian Associates to supply two S-band klystrons, with spares, to the Laboratory. Butman and Gordon L. Guernsey used these klystrons in an S-band amplifier chain to generate 1 MW of peak power and 2 kW of average power. Their driving oscillator was a low-power klystron whose frequency was stabilized by a very high-Q cavity. The klystrons were thoroughly tested, and they found extensive use in an experimental radar at the Laboratory, where

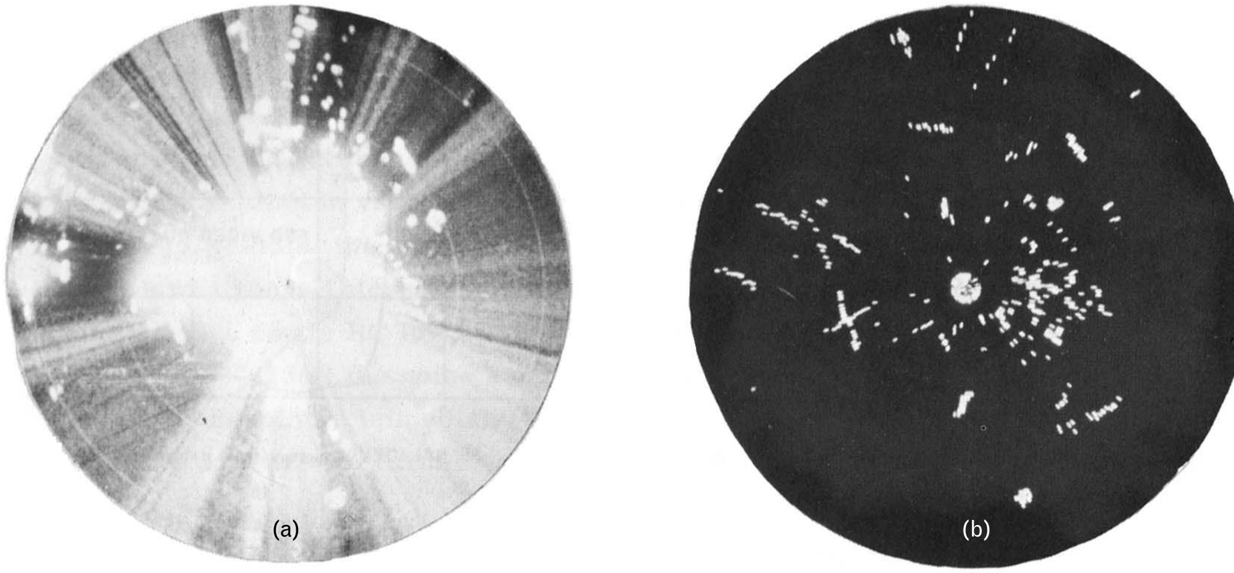

FIGURE 6. Comparison of displays from two receivers subjected to three jammers on a single B-47 at 35,000-ft altitude and 10-mi range: (a) time exposure of an output display of a linear receiver and (b) time exposure of the output of a hard-limiting receiver. The lines of blips represent aircraft tracks.

they demonstrated stable, reliable, and highly coherent operation—a large improvement over the operation of the magnetron.

Somewhat later Butman and Guernsey tested a six-cavity klystron developed by Varian Associates for Hughes Aircraft Company. For this application, the cavities were stagger-tuned for broad bandwidth at the expense of gain. In the Lincoln Laboratory tests the six cavities were synchronously tuned to maximize the gain. A stable gain of 89 dB was attained with root-mean-square phase fluctuation under 2.5°—a remarkable result at that time.

These tests marked the beginning of what came to be an extensive program in which Butman, Guernsey, and Clarence W. Jones worked with industry to develop high-power microwave components and to test numerous high-power klystrons for a variety of applications. The development of klystrons opened up the possibility, exploited in a later era, of using a variety of radar waveforms on demand, such as frequency-modulated pulses for pulse compression and short bursts of closely spaced pulses to permit measurement of very high Doppler frequencies.

Demonstration of an Integrated Radar System

To cap off the work described above we combined several techniques into an integrated demonstration system. The key idea for the system, proposed by Martin Axelbank, was an unusual form of signal integration. He observed that the hard-limited echo from an extended target, such as chaff or precipitation, is decorrelated (i.e., becomes noiselike) if the radar frequency jumps at random from pulse to pulse. By contrast, the echo from a large target, such as an aircraft, has a component that remains steady from pulse to pulse. Integration, necessarily noncoherent, could enhance the signal of an aircraft relative to that of competing chaff or precipitation, thereby yielding a technique later called superclutter visibility.

The system used a two-pulse delay-line MTI, a quartz delay-line analog integrator, and a transmitter that randomly jumped frequency over a broad band of frequencies after each pair of pulses. The combination of the hard-limiting receiver and frequency jumping made the radar highly resistant to jamming. The two-pulse MTI used a short interpulse period that broadened the filter null around zero Doppler frequency and moved the first blind speed out to a high value, which provided reasonably good Doppler filtering. The analog integrator in combination with the frequency jumping was effective in filtering out chaff and rain clutter. This system was as successful as any that we tried. We could not test it at length, however, because other radar operators nearby were distinctly unenthusiastic about the interference effects on their radars of its frequency-jumping mode.

Error Correction

Finally, attention should be called to some work on error correction, somewhat out of the main stream of radar research. As pointed out above, the Cape Cod System transmitted digitized radar data over telephone lines. At that time, digitized data transmission was a largely undeveloped art. Telephone lines that were adequate for voice signals were often far less than ideal for digital signals, resulting in significant errors in the received signals. This inadequacy motivated an investigation of how to reduce errors.

Irving S. Reed and Gustave Solomon investigated a number of mathematical schemes with potential for error correction. Their work culminated in publication of a fundamental and rather abstract mathematical paper in the *Journal of the Society for Industrial and Applied Mathematics* [20]. Although their paper was little noticed at the time, it contained basic ideas that have since been developed into powerful and widely used error-correction schemes, now known as Reed-Solomon error-correcting codes [21]. The codes have been used with compact discs, digital audio tape, high-definition TV systems, and the *Voyager* and *Galileo* spacecraft. Reed and Solomon received the 1995 IEEE Masaru Ibuka Consumer Electronics Award for this work.

Retrospective

As the digital-computer technology needed for the SAGE system matured, the MITRE Corporation was set up in 1958 to oversee the implementation of the system in the field. Related work at Lincoln Laboratory was attenuated. The imminent advent of intercontinental ballistic missiles and the launch of *Sputnik I* further dampened interest in research on air-defense techniques. The work on radar from the Laboratory's first several years, described above, did not have the practical impact on air-defense radars that it might have had otherwise. One fruitful result, however, was the education of a cadre of people at Lincoln Laboratory in both the underlying theory of radar detection and parameter estimation and in the nitty-gritty of practical applications, preparing them to take on new and greater challenges.

Acknowledgments

I am indebted to the caretakers of the Lincoln Laboratory archives for help in finding early Lincoln Laboratory reports and copying pertinent excerpts.

REFERENCES

1. A.G. Hill, J.W. Forrester, and G.E. Valley, "Quarterly Progress Report, Division 2—Aircraft Control and Warning, Division 6—Digital Computer," 1 June 1952. (This report was the first in a long series of quarterly progress reports. Some information from later reports in the series has also been used here.)
2. E.C. Freeman, ed., *MIT Lincoln Laboratory: Technology in the National Interest* (Lexington, Mass., 1995), pp. 15–33.
3. J.V. Harrington, "An Analysis of the Target Detection of Repeated Signals in Noise by Binary Integration," *IRE. Trans. Inf. Theory* **1** (1), 1955, pp. 1–9.
4. W. L. Root, "Remarks, Mostly Historical, on Signal Detection and Signal Parameter Estimation," *Proc. IEEE* **75** (11), 1987, pp. 1446–1457
5. W.B. Davenport, Jr., and W.L. Root, *An Introduction to the Theory of Random Signals and Noise* (McGraw-Hill, New York, 1958).
6. R. Payne-Scott, "The Visibility of Small Echoes on PPI Displays," *Proc. IRE* **36** (2), 1948, pp. 180–196.
7. J.I. Marcum, "A Statistical Theory of Target Detection by Pulsed Radar, and Mathematical Appendix," *IRE Trans. Inf. Theory* **6** (2), 1960, pp. 59–267 (originally published as RAND Corp. Res. Mem. RM-754, 1 Dec. 1947, and RM-753, 1 July 1948).
8. P. Swerling, "MTI Radar," in *Radar Handbook,* M.I. Skolnik, ed. (McGraw-Hill, New York, 1970), pp. 4.1–4.14.
9. E.J. Kelly, I.S. Reed, and W.L. Root, "The Detection of Radar Echoes in Noise," *SIAM J.* **8** (2), 1960, pp. 309–341, and *SIAM J.* **8** (3), 1960, pp. 481–507.
10. E.J. Kelly, "The Radar Measurement of Range, Velocity and Acceleration," *IRE Trans. Mil. Electron.* **5** (2), 1961, pp. 51–57.
11. I.S. Reed, "On the Use of Laguerre Polynomials in Treating the Envelope and Phase Components of Narrow-Band Gaussian Noise," *IRE Trans. Info. Theory* **5**, Sept. 1959, pp. 102–105.
12. R.P. Wishner, "Distribution of the Normalized Periodogram Detector," *IRE Trans. Info. Theory* **8**, Oct. 1962, pp. 342–349.
13. I.I. Shapiro, *The Prediction of Ballistic Missile Trajectories from Radar Observations* (McGraw-Hill. New York, 1958).
14. W.W. Schrader, "MTI Radar," in *Radar Handbook,* pp. 17.1–17.60.
15. J.P. Eckert, Jr., "A Survey of Digital Computer Memory Systems," *Proc. IRE.* **41**, Oct. 1953; reprinted in *Proc. IEEE* **85** (1), 1997, pp. 184–197.
16. W.R. Bennett, "Methods of Solving Noise Problems," *Proc. IRE* **44** (5), 1956, pp. 609–638.
17. D.R. Griffin, J.J.G. McCue, and A.D. Grinnell, "The Resistance of Bats to Jamming," *J. Exp. Zool.* **152**, 1963, pp. 229–250.
18. D.A. Cahlander, J.J.G. McCue, and F.A. Webster, "The Determination of Distance by Echolocating Bats," *Nature* **201**, 8 Feb. 1964, pp. 544–546.
19. J.J.G. McCue, "Aural Pulse Compression by Bats and Humans," *J. Acoust. Soc. Am*. **40** (3), 1966, pp. 545–548.
20. I.S. Reed and G. Solomon, "Polynomial Codes over Certain Finite Fields," *SIAM J.* **8** (2), 1960, pp. 300–304.
21. B.A. Cipra, "The Ubiquitous Reed-Solomon Codes," *SIAM News* **26** (1), 1993.

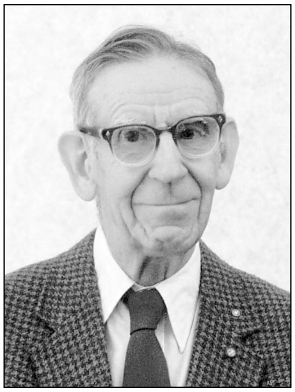

DONALD L. CLARK
received a bachelor's degree in electrical engineering at the University of Vermont in 1943. He worked in the research laboratory at Stromberg Carlson in Rochester, New York, mainly on magnetic recording. He earned a Ph.D. degree in physics at the University of Rochester, where he experimented with pi mesons. He joined Project Lincoln in 1952 and enjoyed a 31-year career at Lincoln Laboratory. He worked mostly on various problems of radars at the cutting edge of technology. As a result of the experience described in the article he was an early advocate for taking into account during the specification and design phase of military radar developments the likely countermeasures to which the radars might be subjected. His group studied test-range measurement radars, ballistic-missile-defense radars and other sensors in a variety of contexts, and satellite surveillance and identification radars, exploiting—where possible—data from the Laboratory's extensive field-measurements program at Kwajalein. His proudest and most lasting accomplishment was bringing a number of outstanding people to the Laboratory and helping to guide their careers. He is a life member of AAAS and IEEE, a senior member of APS, an emeritus member of AAPT, and a member of Sigma Xi.