Radar Signal Processing

Robert J. Purdy, Peter E. Blankenship, Charles Edward Muehe, Charles M. Rader, Ernest Stern, and Richard C. Williamson

This article recounts the development of radar signal processing at Lincoln Laboratory. The Laboratory's significant efforts in this field were initially driven by the need to provide detected and processed signals for air and ballistic missile defense systems. The first processing work was on the Semi-Automatic Ground Environment (SAGE) air-defense system, which led to algorithms and techniques for detection of aircraft in the presence of clutter. This work was quickly followed by processing efforts in ballistic missile defense, first in surface-acoustic-wave technology, in concurrence with the initiation of radar measurements at the Kwajalein Missile Range, and then by exploitation of the newly evolving technology of digital signal processing, which led to important contributions for ballistic missile defense and Federal Aviation Administration applications. More recently, the Laboratory has pursued the computationally challenging application of adaptive processing for the suppression of jamming and clutter signals. This article discusses several important programs in these areas.

The history of radar signal processing at Lincoln Laboratory had its genesis in research efforts undertaken at the MIT Radiation Laboratory during World War II [1]. These efforts, along with similar efforts at Bell Telephone Laboratories [2, 3], provided a theoretical foundation for many important developments in signal processing at many organizations during the ensuing years [4]. With the formation of Lincoln Laboratory in 1951, this theoretical foundation was initially applied to programs in air defense. Soon, however, the stringent needs of ballistic missile defense required the application of both signal processing theory and practice. Subsequently, signal processing requirements from fields as diverse as air traffic control, space surveillance, and tactical battlefield surveillance also stimulated the development and implementation of powerful new signal processing techniques and technology.

The essence of signal processing is its combination of theory, efficient computational algorithms, and the implementation of these algorithms in hardware. One interesting aspect of the history of radar signal processing at Lincoln Laboratory is the transference of techniques developed for one mission area to other mission areas. For example, Lincoln Laboratory's efforts on air defense were applied to the needs of air traffic control, satellite communication contributed to developments in space surveillance, and speech processing and solid state physics both contributed significantly to radar signal filtering. Particularly significant have been the pathfinding efforts in digital signal processing, and the successful application of this field to many important problems across various areas of application.

The SAGE Air-Defense System

In the early 1950s Lincoln Laboratory participated in the first application of digital technology to radar signal processing. The Semi-Automatic Ground Environment (SAGE) Air Defense System was under development, and there was a need to transmit target information from the radars over narrow-bandwidth telephone lines to the direction centers. The solution to this problem was the sliding-window detector illustrated in Figure 1. The name *sliding window* refers to the short length of time that a rotating antenna's

beam dwells upon a target. Each implemented range gate was assigned an accumulator. In each range gate the video output from each radar pulse was sampled and subjected to an initial threshold. This output was assigned a "1" value and added to the accumulator if the initial threshold was exceeded. A "0" value meant no detection and "1" was subtracted from the accumulator. The accumulator was never allowed to go below zero. A target was declared when the sum in the accumulator exceeded a second threshold, as shown in Figure 1(b), and the end of the run was declared when the sum in the accumulator fell below a third threshold. The midpoint between these declarations was generally used as the azimuth estimate of the target, as shown in Figure 1(c). In the absence of a target the receiver noise would normally cause the accumulator sum to hover well below the second threshold. The sliding-window detector approximated what a human operator would do in deciding on the presence of a target on a radar plan-position-indicator (PPI) display, and produced approximately the same noncoherent integration gain as does the human operator. For each detection a single digital word containing range, azimuth, and strength of target was assembled and sent over the telephone line. Analyses of the performance of the sliding-window detector were reported by Gerald P. Dinneen and Irving S. Reed [6]. The sliding-window detector, which was later renamed the common digitizer, became the standard method for detection in long-range ground-based surveillance radars for both air traffic control and military applications.

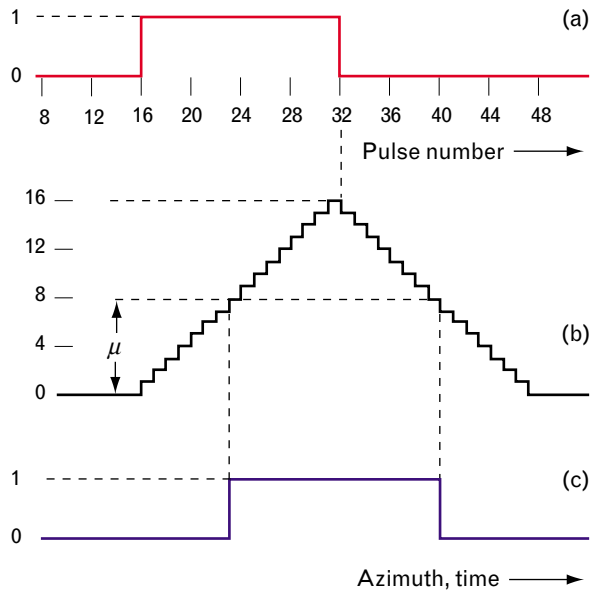

FIGURE 1. The sliding-window detector, operating with ideal signal input. (a) The binary-quantized video signal after the application of an initial threshold to the range gate of interest. (b) The accumulation of the binary count of successive returns from the range gate during one radar-beamwidth traversal time. (c) The resulting binary sequence showing detection of the target when the count exceeds another threshold μ. The beam-split estimate of azimuthal position corresponds to the midpoint of the interval during which the cumulative sum exceeds the threshold μ [5].

Ballistic Missile Defense

With the increasing ballistic missile threat in the 1950s, the Laboratory became heavily involved in developing signal processing technology to address the increasingly sophisticated radar signals that were used to make measurements on ballistic missile reentry complexes. The theoretical basis for radar signal design was advanced by the application of radar ambiguity-function analysis, especially in high-clutter environments [7, 8]. The problem then became one of identifying the appropriate technology for hardware implementation. Initial efforts used commercially available technology and were of limited capability [9]. Fortunately, technology was advancing, and two application areas that were unique to the Laboratory proved to be particularly successful: surface-acoustic-wave signal processing and digital signal processing.

Surface-Acoustic-Wave Signal Processing

In the late 1960s a number of researchers around the world became interested in the potential use of surface acoustic waves (SAW) for providing new types of compact filters that could operate in a frequency range from a few tens of hertz up to a few gigahertz. Among other applications, the projected device parameters seemed well matched to implementing analog pulse-compression filters for radars. As a result, the development of SAW devices for military use began in several laboratories. One of the earliest efforts was established at Lincoln Laboratory under the leadership of Ernest Stern [10–12]. In the late 1960s this group began to pursue the development of SAW devices for radar and communications applications.

The First Reflective Array Compressors

The challenge for SAW technology was to achieve sufficiently precise devices with the right combinations of correlation time and bandwidth to be useful in radar systems. The earliest SAW dispersive delay lines for use as radar-pulse compressors employed a metallic pattern of interdigitated electrodes deposited on the surface of a piezoelectric crystal such as lithium niobate [13, 14]. The electrodes launched an acoustic wave on the crystal surface; the electrode pattern was arranged so that it would be responsive to the specific received signal. This interaction yielded the desired chirp response. The approach worked reasonably well for bandwidths of a few tens of megahertz and for time-bandwidth products of one hundred or less, but it failed to yield sufficiently precise response and low sidelobes when tried at higher time-bandwidth products.

Results obtained previously with some low-bandwidth acoustic filters [15] suggested to the Lincoln Laboratory SAW group that reflection of SAWs from arrays of grooves etched into the crystal surface could yield a more nearly ideal device response than that obtained with metallic electrode arrays. To explore this hypothesis, the SAW group realized that experiments were needed to elucidate the physics of surface wave reflections, new technology was needed to lithographically define and etch the reflective arrays, and new device models and design techniques had to be developed.

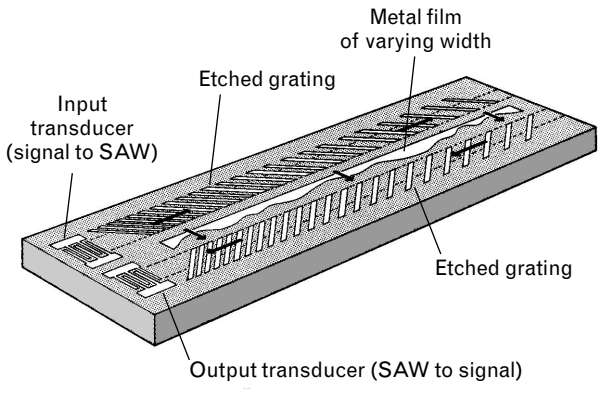

FIGURE 2. A phase-compensated reflective-array compressor, or RAC. The input transducer converts an electrical signal into a surface acoustic wave (SAW) that propagates along the surface of the crystal. The grating etched into the crystal reflects the wave at a position determined by the input frequency and the local spacing of the grooves in the grating. High frequencies reflect close to the input transducer, while low frequencies reflect at the far end of the grating. A second reflection sends the SAW to the output transducer, where it is converted back into an electrical signal. The desired delay versus frequency is set by the geometry of the device. Deviations from the desired response can be trimmed out by a metal film of varying width deposited on the device.

A great deal of the technological groundwork for this process was established during 1971. By 1972, fabrication of the first reflective-array compressor (RAC) was initiated; this device is illustrated in Figure 2. The first RAC device was a linear-FM filter with a 50-MHz bandwidth (on a 200-MHz carrier) matched to a 30-μsec-long waveform [16–18]. This arrangement yielded a time-bandwidth product of 1500, more than an order of magnitude greater than that achieved by interdigital-electrode SAW devices [19]. The response was remarkably precise; the phase deviation from an ideal linear-FM response was only about 3° root mean square (rms). Pairs of matched RACs were used in pulse-compression tests in which the first device functioned as a pulse expander and the second as a pulse compressor. The compressed pulsewidths and sidelobe levels were near ideal. Armed with these encouraging results, researchers took the next step by developing RAC devices for specific Lincoln Laboratory radars.

RAC Pulse Compressors for the ALCOR Radar

The ARPA-Lincoln C-band Observables Radar, or ALCOR [20], on Roi-Namur, Kwajalein Atoll, Marshall Islands, had a wideband (512 MHz) 10-μsec-long linear-FM transmitted-pulse waveform (see the article entitled "Wideband Radar for Ballistic Missile Defense and Range-Doppler Imaging of Satellites," by William W. Camp et al., in this issue). ALCOR was a key tool in developing discrimination techniques for ballistic missile defense. The wide bandwidth yielded a range resolution that could resolve individual scatterers on reentering warhead-like objects. This waveform was normally processed with the STRETCH technique, which is a clever time-bandwidth exchange process developed by the Airborne Instrument Laboratory [21, 22]. The return signal is

mixed with a linear-FM chirp and the low-frequency sideband is Fourier transformed to yield range information. For a variety of reasons, the output bandwidth and consequently the range window were limited. For example, the ALCOR STRETCH processor yielded only a thirty-meter data window. Therefore, examination of a number of reentry objects, or the long ionized trails or wakes behind some objects, required a sequence of transmissions.

This sequential approach was inadequate in dealing with the challenging discrimination tasks posed by reentry complexes, which consist not only of the reentry vehicle, but also a large number of other objects, including tank debris and decoys, spread out over an extended range interval. What was needed was a signal processor capable of performing pulse compression over a large range interval on each pulse. Lincoln Laboratory contracted with Hazeltine Laboratory to develop a 512-MHz-bandwidth all-range analog pulse compressor employing thirty-two parallel narrowband dispersive bridged-T networks built out of lumped components, to cover the bandwidth. The resulting processing unit, shown in Figure 3, was large (it filled about seven relay racks) and complex, and it required a great deal of tweaking to yield reasonable sidelobes. Cost and complexity loomed large when plans were made for a series of reentry tests in which matched pairs of pulse compressors would be required. In a parallel effort, the Lincoln Laboratory SAW device group was challenged to develop pulse compressors that could meet the all-range needs of ALCOR. This task would mean extending the band-

FIGURE 3. The ALCOR all-range wideband analog pulse compressor developed jointly by Lincoln Laboratory and Hazeltine Laboratory.

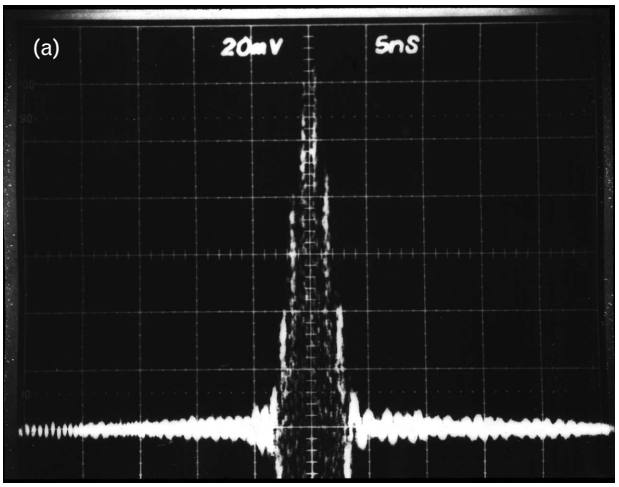

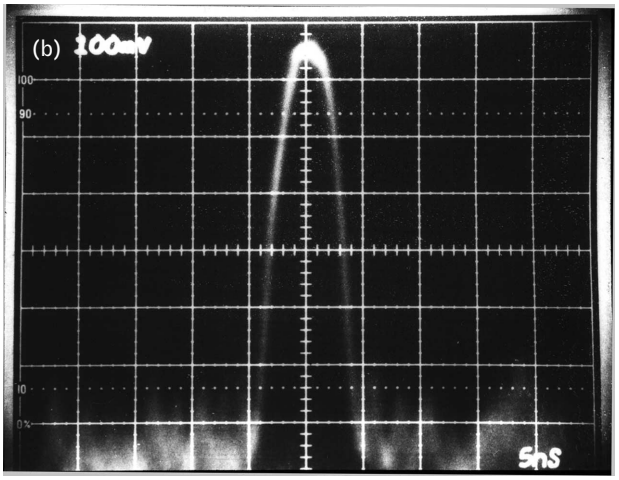

FIGURE 4. RAC sidelobe performance in compressing a 10-μsec 512-MHz-bandwidth pulse. (a) The compressed pulse and its sidelobes on a 1-GHz carrier frequency, shown on a linear scale. (b) The envelope of the compressed pulse and its sidelobes on a logarithmic scale of approximately 6 dB per division. The horizontal scale on both graphs represents 5 nsec per division.

width of SAW RAC technology by an order of magnitude, which would increase the time-bandwidth product well beyond that achievable with any existing analog device technology.

During 1972 and 1973, Lincoln Laboratory developed a 512-MHz-bandwidth (on a 1-GHz intermediate frequency [IF]) 10-μsec RAC linear-FM pulse compressor [23]. In ALCOR, an active circuit with feedback generated the linear-FM chirp, so that the RAC devices were to function as all-range pulse compressors matched to that waveform. To suppress range sidelobes, a Hamming window was built into the RAC devices by varying the etch depth of the grooves as a function of position.

Midway in the development effort, significant difficulty was encountered in achieving sufficiently precise amplitude and phase responses. Subtle lithographic and etching effects yielded errors in groove depths and positions that measured only a few tenths of a nanometer, but these very small errors were large enough to degrade the compressed-pulse sidelobes significantly. A trimming technique was developed to achieve an adequately precise response. This technique required measuring the device and the subsequent deposition of a corrective metal pattern of varying width on the crystal surface of the RAC, as illustrated in Figure 2. The resulting precision allowed for a phase response that was precise to about 2.5° rms, or about one part per million over the 5120 cycles of the waveform. This response yielded near-in-range sidelobes in the –35-dB range, whereas far-out sidelobes rapidly fell to better than 40 to 50 dB down, as shown in Figure 4. In Figure 5, which is a photograph of a RAC developed for ALCOR, the two rainbow-colored stripes near the centerline of the crystal show light that is diffracted from the etched grating. The phase-compensating varying-width metal film strip runs down the centerline of the crystal.

Pairs of approximately one-inch-long matched RAC devices were installed in ALCOR in 1974 and were used successfully in a series of reentry tests. These devices proved to be such powerful wide-bandwidth signal processors that advances in analog-to-digital converter technology to capture the output were required before the capability of the RAC devices could be fully utilized.

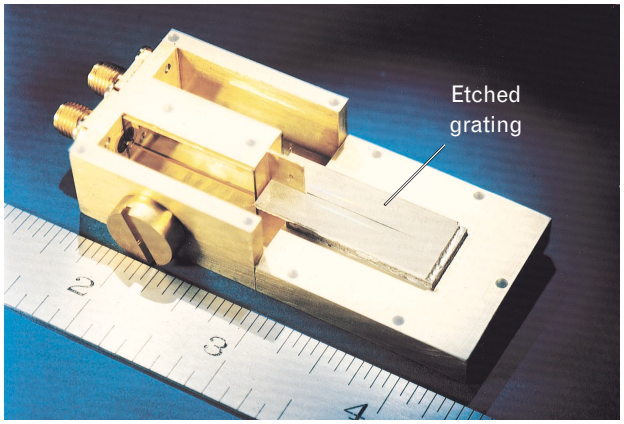

FIGURE 5. The ALCOR RAC processor. The two rainbow-colored stripes near the centerline of this device are created by the diffraction of light off the pair of etched gratings. The varying-width metal film strip running along the centerline of the device performs phase compensation. This device replaced the entire seven-rack processor shown in Figure 3.

RAC Pulse Compressors for the MASR Airborne Radar

Following the positive results with the early RAC devices, SAW technology was considered for a number of Lincoln Laboratory programs. As the technology matured, the Laboratory SAW group helped guide the development and procurement of SAW devices from outside companies. Some device specifications fell outside the state of the art, however, and so the initial development of these more challenging devices was carried out at the Laboratory. One example was the pulse compressors required for the experimental Multiple-Antenna Surveillance Radar (MASR), an airborne radar for ground surveillance (see the article entitled "Displaced-Phase-Center Antenna Technique," by Charles Edward Muehe and Melvin Labitt, in this issue). This radar employed a 2.5-MHz-bandwidth pulsed linear-FM waveform with a duration of 125 μsec.

The long pulse in the MASR proved to be a challenge for SAW technology. A new material, bismuth germanium oxide, with a low acoustic velocity was tried. A host of detailed technical obstacles were overcome in order to adapt the RAC technology to this new substrate material [24]. The package developed for MASR incorporated three matched devices: a pulse expander and two weighted pulse compressors. Phase errors were less than 2° rms, yielding better

than –35-dB near-in sidelobes. The RAC devices played an important part in achieving successful detection of slow-moving ground targets from the MASR airborne platform.

Fast Spectrum Analyzers for the Infrared Airborne Radar

Since the 1950s, Lincoln Laboratory and other organizations have realized that sets of dispersive delay lines can be used to implement a high-speed analog Fourier transform by means of the chirp-transform algorithm. The advent of precision SAW dispersive delay lines reawakened this interest in the late 1970s.

The high carrier frequency of the coherent Infrared Airborne Radar (IRAR) provided the unique capability of being able to measure the Doppler shift of target returns with high resolution (approximately one meter per second) in only a few microseconds. Performing the required Fourier transform for incoming target returns in such a short time was very challenging. The task was made even more difficult because the receiver for this CO_2 laser-radar system employed a twelve-element array of photomixers, thus requiring that spectral analysis be performed on twelve parallel channels simultaneously. A compact processor consisting of twelve RAC-based chirp-transform units was developed to accomplish this task. Figure 11 in the article in this issue entitled "Development of Coherent Laser Radar at Lincoln Laboratory," by Alfred B. Gschwendtner and William E. Keicher, shows typical results achieved with this system

Memory Correlators

Whereas reflector gratings are fixed matched filters, a major effort was invested in realizing programmable devices capable of responding to a variety of waveforms. Chief among these devices are acousto-electric convolvers [25, 26] that act as matched filters to continuously changing waveforms for spread-spectrum communication equipment, such as DARPA's packet radio program. These devices achieved bandwidths of 100 MHz and duration times of 10 μsec or more. A programmable matched filter, called a memory correlator, was invented and developed for use in advanced radar demonstrations, with similar bandwidths and time-bandwidth products [27, 28].

The Legacy of the SAW Development Effort

The Communications division at the Laboratory at that time was formulating plans for a new satellite communications system that would have increased jamming resistance and capacity for simultaneous multiple access by many authorized users. Fast frequency hopping [29] had a clear advantage for jamming resistance, since fast hopping implied a short dwell time on each frequency. The short dwell time required that the receiver circuitry demodulate the information in a time period that was too short for the digital circuitry of the era to accommodate. A fast SAW spectrum analyzer was developed to meet these requirements [30], and it was incorporated into two Fleet Satellite (FLTSAT) extremely high frequency (EHF) packages (FEP), which were launched as extra features of the satellites FLTSAT-7 and -8, launched in 1986 and 1989, respectively. Each FEP contains five SAW devices; they have functioned flawlessly in orbit since launch.

The Laboratory's demonstration that SAW reflection gratings could yield precision device response in matched filters stimulated an interest at a number of laboratories in applying grating technology to other purposes. A key advance was the demonstration that a high-performance resonator could be made [31]. This development in turn led to two areas of significant application: low-noise oscillators and narrowband filters for commercial and military equipment.

The conventional lithographic fabrication technology available in the 1970s was not capable of producing the precise high-resolution large-area patterns required for SAW devices. For the high-frequency devices, the lines in the SAW transducers and the reflection gratings were less than a micron wide, well beyond the state of the art at that time. As a result, considerable effort was spent developing advanced techniques such as improved pattern generators, electron-beam lithography, and advanced photoresist procedures. The Lincoln Laboratory SAW group invented X-ray lithography as a means to reproduce fine lithographic features [32]. Many elements of advanced lithography were applicable to a wide range of devices, not just to SAW devices, and the lithography effort took on a life of its own. Eventually, a sub-

micrometer technology group was set up at the Laboratory to pursue advanced lithographic techniques. When interest in this area grew on the MIT campus, the Laboratory's expertise was called upon in the establishment of the Microsystems Center at MIT. In addition to transferring lithographic technology, the Laboratory has continued its own role of leadership in microcircuit fabrication techniques.

Digital Signal Processing

The development of digital signal processing for radar at Lincoln Laboratory provides a classic example of interdisciplinary technology transfer. The original efforts of researchers at Bell Telephone Laboratories, and by Bernard Gold and Charles Rader at Lincoln Laboratory [33], were motivated by the desire to bandwidth-compress speech for more efficient digital secure-voice communication and to digitally simulate analog components. This work led to Gold and Rader's seminal book on digital signal processing [34]. The techniques developed during this time were very powerful, and their immense applicability to signal processing for ballistic missile defense became readily apparent [35].

The key realization of the potential for digital signal processing in radars was the understanding that ballistic-missile-defense radars are pulsed systems and, unlike analog signal processing, the digital signal processing did not need to be time synchronous. If raw data are digitized [36] and stored in memory, the available processing time is the time until the next measurement, not the real-time extent of the measurement itself. This approach, then as it is now, is a careful balance among the required algorithms, an architecture that efficiently but flexibly implements those algorithms, and the selection of a hardware technology that meets timeline requirements. Examples of both the programmable and special-purpose approaches to radar signal processing are described below.

The Fast Digital Processor

In the mid-1960s the emergent field of digital signal processing was becoming more well known. Exciting new techniques for designing and implementing digital filters were being published, and the fast-Fourier-transform (FFT) algorithm in its various incarnations offered the prospect of drastically reducing the number of computations necessary to perform important signal processing functions digitally (primarily multiplications, which were time-consuming operations on a general-purpose computer).

At Lincoln Laboratory there was growing frustration among researchers over the inadequacy of the general-purpose computer technology of the day for performing digital-signal-processing calculations with any kind of reasonable speed, notwithstanding computationally efficient algorithms such as the FFT. Thus in 1967, a team led by Gold, Rader, and Paul McHugh conceived the architecture and instruction set for the Fast Digital Processor (FDP) [37, 38]. Although, as mentioned above, a driving motivation was to simulate developmental speech-coding algorithms in real time or near real time, the overarching goal of the project was to achieve a design representing an optimum balance for digital-signal-processing applications between the computation throughput-rate potential offered by a purely special-purpose architecture and the flexibility afforded by a general-purpose computer. The result was a programmable machine, architecturally optimized for digital-signal-processing computations, that offered the prospect of approximately two hundred times the throughput rate of a general-purpose computer for many digital-signal-processing applications through a combination of advanced digital integrated-circuit technology (emitter-coupled logic), architectural parallelism, instruction pipelining, and clever specialized architectural features (e.g., a "bit-reversed add" to facilitate radix-2 FFT address calculations).

The FDP architecture, illustrated in Figure 6, used distinct structures for the program and data memories, and it used a semimicrocoded instruction set. The FDP featured a 512 × 36-bit program memory to support the wide instruction-word format, which was physically separate and distinct from two simultaneously accessible 1024 × 18-bit data memories (extendable to 4096), all of which were implemented with semiconductor-memory technology. The architecture also incorporated four identical 18-bit, twos-complement, fixed-point arithmetic elements, as illustrated in Figure 7, which could be operated

concurrently and independently by virtue of the latitude provided by the 36-bit-wide instruction word. The FDP designers were among the first in the field of digital signal processing to recognize the so-called multiply-accumulate operation as the most elemental digital-signal-processing computational building block, and the complex multiply as fundamental to FFT calculations. Therefore, the arithmetic elements were configured and interconnected to facilitate these critical types of operations. The FDP was also equipped with flexible and powerful data-memory address-calculation mechanisms to further enhance efficiency and performance for a wide class of digital-signal-processing functions.

The timing of the FDP was based on a three-deep instruction pipeline comprising three 150-nanosecond epochs, which overlapped instruction fetch with instruction decode/data-memory access and arithmetic-element operations. In principle, it was possible to perform four arithmetic operations and four local-data transfers per 150-nanosecond epoch, representing a peak theoretical throughput rate of approximately 53 million instructions per second (MIPS). The four-quadrant multiplier, the single most costly component in the arithmetic elements in terms of hardware complexity, was implemented as fully instantiated combinatorial-logic arrays based on a modified Booth's algorithm, and required 450 nanoseconds to produce a signed 36-bit product. To mitigate this extra delay, other operations could be conducted within an arithmetic element while a multiplication was in process.

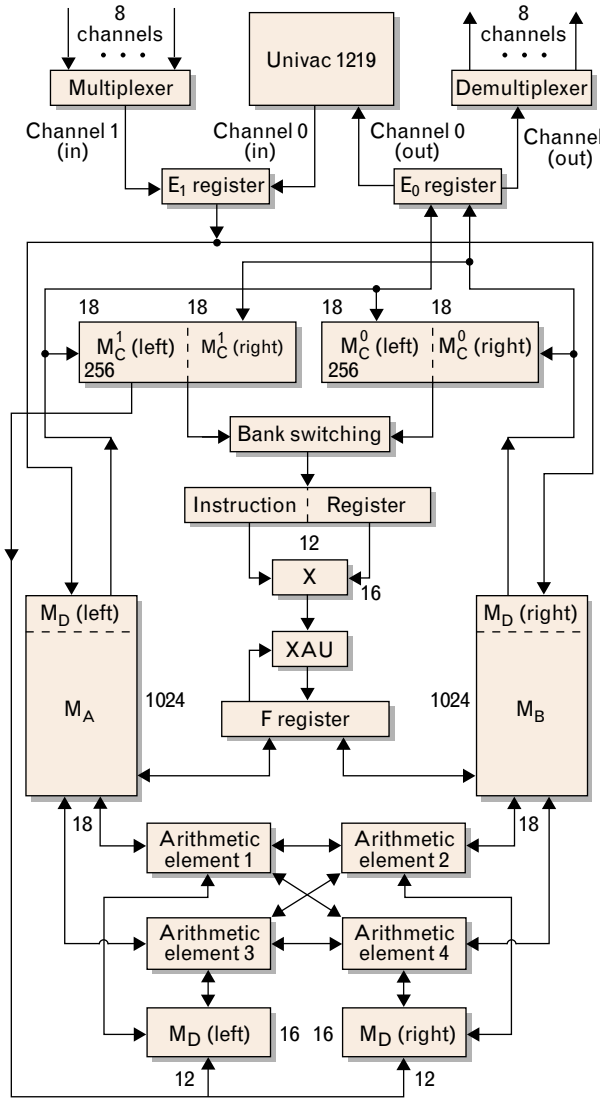

FIGURE 6. The Fast Digital Processor (FDP) architecture. The FDP itself comprised approximately 15,000 emitter-coupled-logic integrated circuits, dissipated about 2.5 kilowatts of power, and occupied about 200 cubic feet of volume. As technology evolved, an equivalent amount of computing power could be realized in a few cubic feet. Such machines were known as *array processors*.

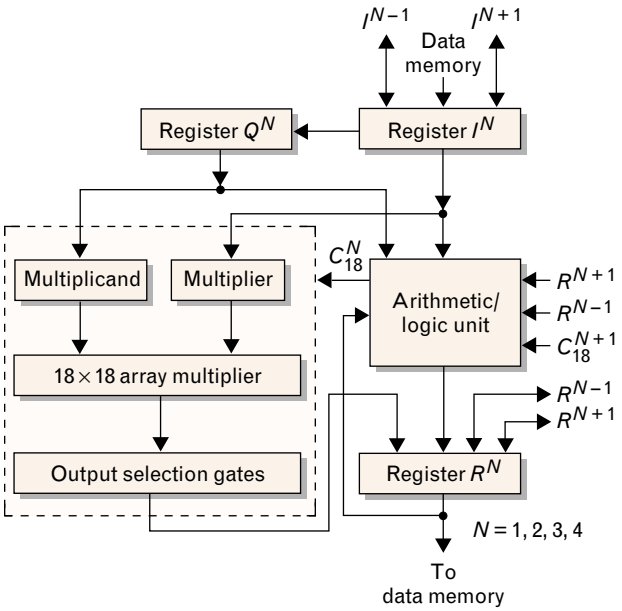

FIGURE 7. FDP arithmetic-element structure. The design showed parallelism in several forms, including dual data memories, four identical arithmetic elements, and a separate program memory. These features provided enhanced performance, particularly when computing complex arithmetic.

FIGURE 8. The FDP facility at Lincoln Laboratory in 1970, which included a Univac 1219 general-purpose host computer. The arithmetic/logic unit incorporated a full 18-bit, twos-complement adder/subtractor, supported all Boolean functions, and included linkages for extended-precision calculations. The 18 × 18-bit four-quadrant multiplier was based on a modified Booth's algorithm, and was implemented as a full combinatorial array using single-bit adders.

Actual design and fabrication of the FDP were carried out at Lincoln Laboratory during the time frame from 1968 to 1970, and represented no mean engineering feat. Some of the innovative layout and packaging concepts incorporated in the FDP came from the people in the Engineering division who had been building the Lincoln Experimental Satellites (LES). To achieve the desired performance goals for the FDP, the design and fabrication team needed to capitalize on the then state-of-the-art Motorola MECL II small-scale and 10k medium-scale digital integrated-circuit technologies. This effort required the development of novel and sophisticated design methodologies heretofore unheard of in digital system implementations, because of the high speed of the logic and the finite speed of electrical-signal propagation. For example, all data, control-signal, and clock-distribution paths required careful attention to physical length, signal quality, and impedance control for reliable and predictable operation. The design practices pioneered in the construction of the FDP eventually became commonplace within the digital design community as experience with ultrahigh-speed digital-circuit technology grew. Figure 8 shows the finished FDP facility, which included a Univac 1219 general-purpose host computer. The FDP proper comprised approximately 15,000 integrated circuits, dissipated about 2.5 kilowatts of power, and occupied nominally 200 cubic feet of space.

Although not easy to program, the FDP proved to be a unique, versatile, and powerful asset, as had been hoped. For example, a two-pole digital resonator or a radix 2 FFT "butterfly" could be executed in approximately 1.2 μsec. The architecture, though optimized for digital filtering and FFT computations, was still general enough to be useful for other types of numeric computation, and it even supported extended-precision and floating-point operations. As a testi-

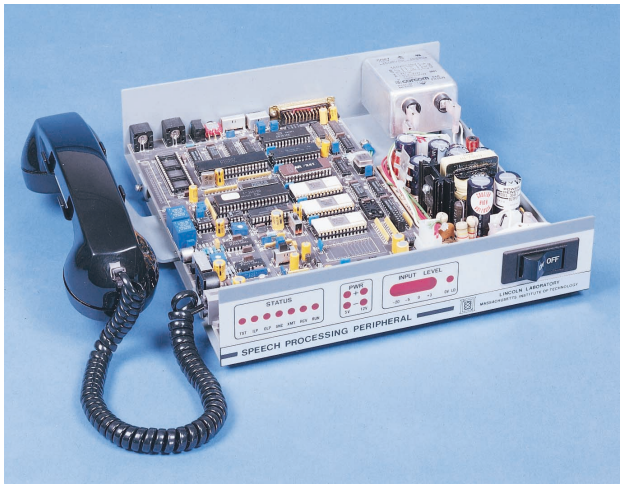

FIGURE 9. The first stand-alone compact linear predictive speech coder, or LPC vocoder, which served as a major driver and motivating force for the next-generation commercial secure telephone units (STU-III) introduced into government service during the early 1980s. This vocoder was based on the state-of-the-art commercial single-chip digital-signal-processing microcomputers available at that time.

mony to its flexibility, the first real-time implementation of a 2400-bit-per-second linear predictive speech coder (LPC), which involved numerical computations far less regular and structured than those of a digital filter or an FFT, was successfully demonstrated on the FDP in the early 1970s [39, 40].

This work led to a series of increasingly compact specialized digital signal processors for real-time implementation of LPC and other digital voice-compression algorithms, such as the first stand-alone LPC vocoder shown in Figure 9. This work culminated in the DARPA-sponsored Speech Processing Peripheral, which was a direct precursor to the next generation of secure telephone units, or STU-IIIs, introduced into government service during the early 1980s.

The FDP also proved useful in radar signal processing applications, where it was capable of real-time performance if appropriate specialized adjunct hardware components were provided when necessary (e.g., an external corner-turning buffer memory) and the range-Doppler space of experimental interest was suitably restricted. For example, in the early 1970s the Federal Aviation Administration (FAA) was exploring signal processing techniques that might provide a cost-effective performance upgrade for the then existing generation of airport surveillance radars (ASR). The FDP was connected to a remote ASR transceiver through a custom-designed duplex data link, and was used for the development and real-time evaluation of novel Doppler-processing techniques for clutter mapping. The FDP simulation experiments proved that a special-purpose digital-signal-processing hardware adjunct to the ASR sensor could be both effective and economical [41, 7]. Also, in a similar time frame, the FDP/data-link facility was used as part of the Long-Range Demonstration Radar project to develop moving-target-indication (MTI) algorithms for surface vehicles or other relatively slow-moving objects amidst heavy ground clutter in defense applications. In particular, these processors allowed engineers to implement and perform real-time evaluations of experimental Doppler-processing, post-detection integration, and statistical-decision algorithms [42].

Although it was a one-of-a-kind machine, the FDP proved the value of programmable machines oriented toward digital signal processing, and it served as a motivator for the first generation of commercial off-the-shelf programmable digital-signal-processing accelerators that reached the marketplace during the 1970s, offered by such manufacturers as Computer Signal Processors, Signal Processing Systems, and Floating Point Systems.

Digital Convolver System

An early application of the special-purpose approach to digital signal processing arose from initial research for the U.S. Army in the early 1970s on an all-solid-state radar for ballistic missile defense, which led to the development of a conceptual L-band radar called the Advanced Fielded Array Radar (AFAR). The L-band radar concept used solid state transmit modules; consequently, it required long waveforms for detection as well as short waveforms for tracking. The need for a large bandwidth to provide adequate range resolution led to a large waveform repertoire with a wide diversity of time-bandwidth products. This repertoire precluded the use of analog filters, in that there simply would have been too many fixed filters.

The flexibility of digital signal processing [43] suggested that a suitable digital processor design would

provide a solution, as long as the processing could be done in real time (i.e., in the total time available). The result was the Digital Convolver System (DCS) [44], which was intended to provide the required flexible real-time matched filtering of large numbers of waveforms, some with large time-bandwidth products. The design was based on fast-convolution techniques [45], and provided for a 16,384-point radix-4 FFT, clocked at 30 MHz, to achieve a throughput data rate of one 16k FFT every 136 microseconds [46].

Two innovations at the time were the use of a hybrid floating-point data format and CORDIC (coordinate-rotation digital computer) [47] rotators in the FFT. The hybrid floating-point format uses a common exponent for both the real and the imaginary parts of the complex data at each stage of the FFT calculation, and it was sometimes referred to as vector floating point. This approach greatly alleviated the computational hardware complexity of the system [48, 49]. Similarly, the CORDIC rotator provided a computationally efficient implementation of the complex multiplications required in the FFT. Another innovation was based on the observation that Doppler processing using the fast-convolution approach did not require the repetitive use of the forward FFT. Rather, a single forward transform followed by multiple inverse transforms was sufficient. The resulting reduction of the hardware requirements (by roughly one half) was significant.

Figure 10 illustrates the DCS architecture. The system includes a temporary storage memory, a reference-function memory, and a multiplier system. The temporary storage memory holds the forward-transformed data and sends the data through the frequency-domain multiplier for multiple inverse transforms. The core of the system is the pipelined FFT [50, 51], which is shown in detail in Figure 11. The most important feature of this system is that the interstage delay-line memories are reconfigurable, which allows the same set of hardware to provide both forward and inverse transforms of 4k, 8k, or 16k points, while also allowing the data to be read into the forward FFT and out of the inverse FFT in normal order. Figure 11 shows seven elementary computation elements and six interstage-delay memory elements, which are reconfigured depending on the size

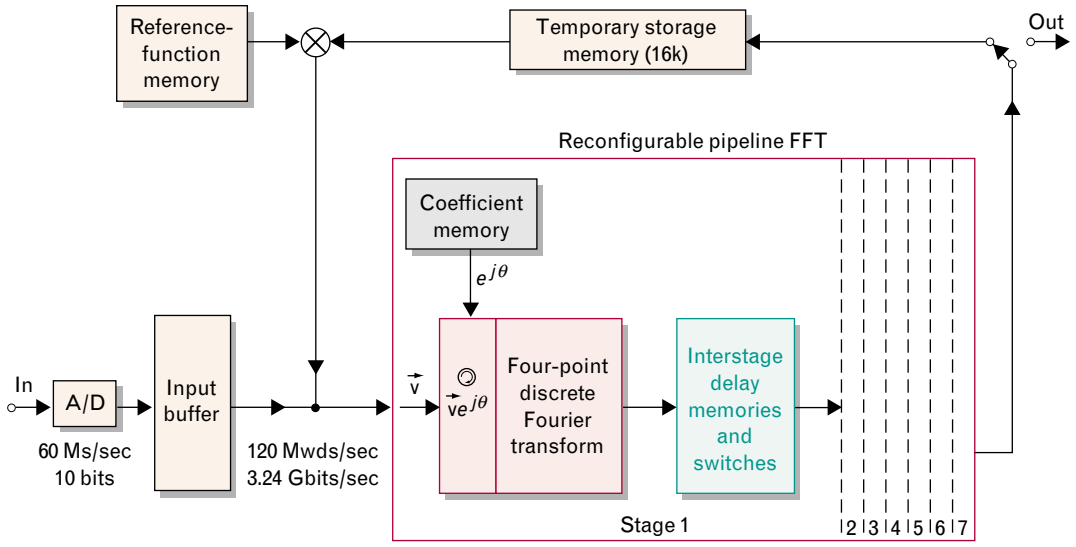

FIGURE 10. The Digital Convolver System (DCS) architecture. This system exploits the fact that Doppler processing of radar waveforms uses Doppler-shifted versions of a single reference function. Consequently, if the processing is performed by fast convolution, only one forward transform is needed. The result is stored, read multiple times, Doppler-shifted, and inverse-transformed multiple times. The forward and inverse transforms are both performed in the reconfigurable pipeline fast-Fourier-transform (FFT) subsystem shown in the figure.

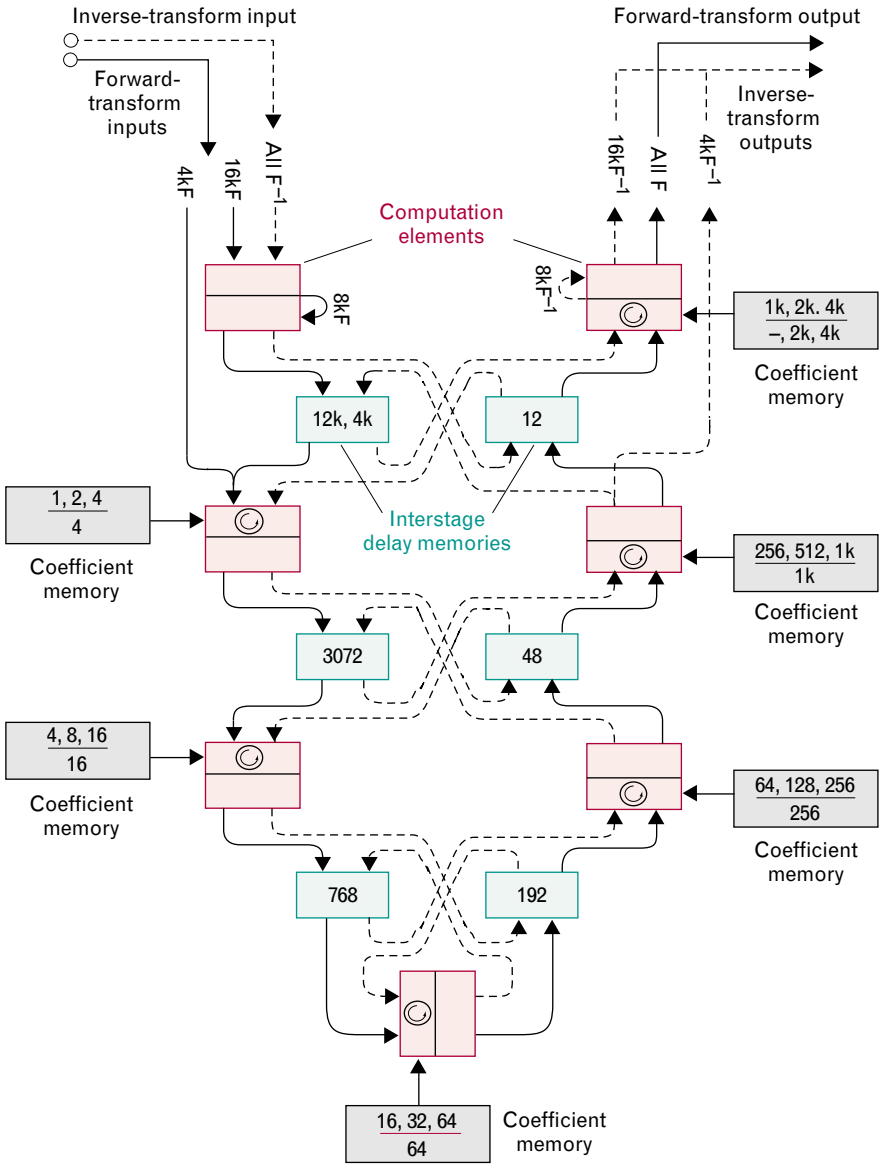

FIGURE 11. The reconfigurable DCS FFT architecture. This system is designed to allow the same hardware subsystems to perform multiple transform sizes (4k, 8k, and 16k) and simultaneously perform both the forward and inverse transforms. The penalty is an increased amount of data routing, but this penalty is more than outweighed by the savings in hardware that would be incurred if two complete transform systems had to be built.

of the transform and whether a forward or inverse transform is being performed. This process is indicated by the two major paths through the figure.

The concept of implementing the signal processing by using digital technology was relatively new at the time. The potential for achieving highly accurate processing, however, was enormous. The DCS demonstrated and certified this potential by achieving a computational noise floor with spurious peaks approximately 63 dB down, as shown in Figure 12, a result that proved the viability of the hybrid floating-point approach.

The DCS [52] used mostly emitter-coupled logic 10k-series integrated circuits to meet the throughput-rate requirements. One large multiplexed memory, however, used MOS technology, and there were a few transistor-transistor-logic interface circuits. The DCS had about 27,500 integrated circuits and consumed

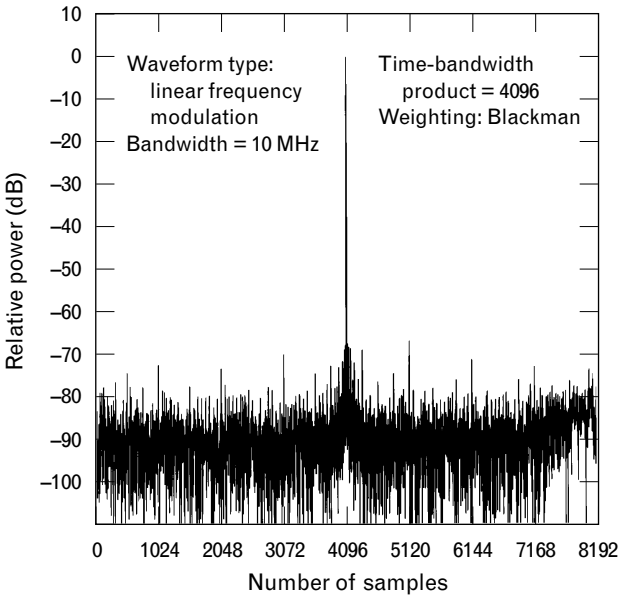

FIGURE 12. The DCS computational noise floor is achieved by using hybrid floating-point arithmetic. The results achieved by the DCS demonstrated that digital-signal-processing techniques have a performance potential limited only by the word length used.

15 kW of power. At its completion in 1979, the DCS, shown in Figure 13, was the fastest and largest pipelined FFT processor that had yet been built.

The FAA: The Moving-Target Detector and the Parallel Microprogrammed Processor

In 1972 the FAA brought a radar problem to Lincoln Laboratory. The FAA was in the process of developing the Automated Radar Terminal System (ARTS-3), with the aim of computerizing air-traffic-control displays at airports. They had successfully automated the Air Traffic Control Radar Beacon System, and in so doing provided automatic track acquisition and updating for all beacon-equipped aircraft (secondary radar). They had been unsuccessful, however, in automating the primary, or skin-tracking, radar. The primary radar produced too many clutter-related false alarms and missed detections as a result of the techniques employed to deal with the clutter.

With the advent of medium-scale integrated circuits around 1970, many new signal processing algo-

FIGURE 13. The DCS in 1979. At that time it was the fastest and largest pipelined FFT processor ever built. The system was large; it was comparable to the ALCOR all-range processor shown in Figure 3, but with an order-of-magnitude improvement in performance.

rithms were developed. This evolving integrated-circuit technology allowed digital sampling and filtering of an ASR's single-scan output in over three million range-azimuth-Doppler cells. Thresholding algorithms (which are described later in this article) could then be employed for the type of clutter found in each resolution cell (i.e., ground clutter in each zero-velocity Doppler cell could be thresholded by using a digitally stored ground-clutter map), thus avoiding false alarms while keeping all of the resolution cells as sensitive as possible for the detection of aircraft. This type of processor was named the moving-target detector (MTD) to distinguish it from the now old-fashioned moving-target indicator (MTI). An initial exercise using Lincoln Laboratory's FDP [37, 41] verified the usefulness of these algorithms over a small eight-nautical-mile by 45° sector. This advance was followed by full-scale development and testing of the MTD, led by Charles Edward Muehe. Two versions of this processor were built. In the MTD-1 the algorithms were hard wired into the processor [53], and in the MTD-2 [54] the algorithms were implemented as software in a parallel microprogrammed processor (PMP) [55]. The MTD-2 found its way into at least six different types of surveillance radars, including both ground-based and airborne radars.

The MTD Class of Radars

The MTD radars incorporate a number of novel signal processing techniques. The older MTI radar's staggered pulse-repetition-frequency (PRF) waveform, which was used to ameliorate blind speeds, is replaced in both the MTD-1 and the MTD-2 by a multiple-PRF waveform, wherein about eight pulses at one PRF in a coherent processing interval are alternated with a coherent processing interval with a 20% different PRF. The receiver maintains linearity over the full dynamic range of the analog-to-digital converters. For each coherent processing interval a bank of digital filters, each designed to maximize the signal-to-clutter ratio, is implemented in each range gate. Several forms of detection thresholding are used, depending on the statistics of the expected clutter reflections in each filter. An algorithm is employed to flag range gates that contain interfering pulses.

To cause a uniform presentation on the plan-position-indicator (PPI) display in the presence of ground clutter, older MTI radars employed amplifiers in the MTI channel that were limited to about 20 dB above the receiver noise [56]. This limiting spreads the clutter spectrum and reduces the MTI subclutter visibility to at most about 20 dB. The MTD, on the other hand, has a measured subclutter visibility of 42 dB, which is in turn limited by the receiver's dynamic range. Because the spatial statistics of ground clutter are highly non-Gaussian, both MTD radars use a clutter map for thresholding the zero-velocity Doppler filter. Older MTI radars have a notch-at-zero Doppler, and thus they cannot detect a crossing target that has a near-zero radial velocity. The clutter map allows detection of crossing aircraft, which would usually present large reflections from their fuselages when crossing or are in a low ground-clutter region because of ground shadowing. As a consequence of this detection capability, the MTD is said to have superclutter and interclutter visibility.

The high pulse-to-pulse correlation of rain-clutter returns, together with noncoherent binary integration, caused the sliding-window detector used in older MTI radars to exhibit a high false-alarm rate in rain. The strictly coherent integration for each of the MTD's nonzero Doppler filters, together with thresholds based on the mean clutter level within ±0.5 nmi of each thresholded range-Doppler cell, keeps the MTD's false-alarm rate under excellent control. The update of the zero-velocity ground-clutter thresholding map is adjusted so that it also keeps up with changing rainstorm backscatter as the storm passes through the radar's coverage. Because multiple PRFs are used, the target appears in a different filter on successive coherent processing intervals (unless it has the same radial velocity as the storm), resulting in a good chance of detection. The MTD's constant PRF in each coherent processing interval, instead of the older MTI radar's staggered pulse-repetition intervals, allows the illumination of second-time-around clutter, which is filtered in the same way as close-in clutter.

For each threshold crossing, a primitive report is sent to the MTD's post-processor, giving the amplitude, range, azimuth, Doppler-filter number, and PRF. Reports that appear to come from the same target are interpolated for the best estimate of the target's

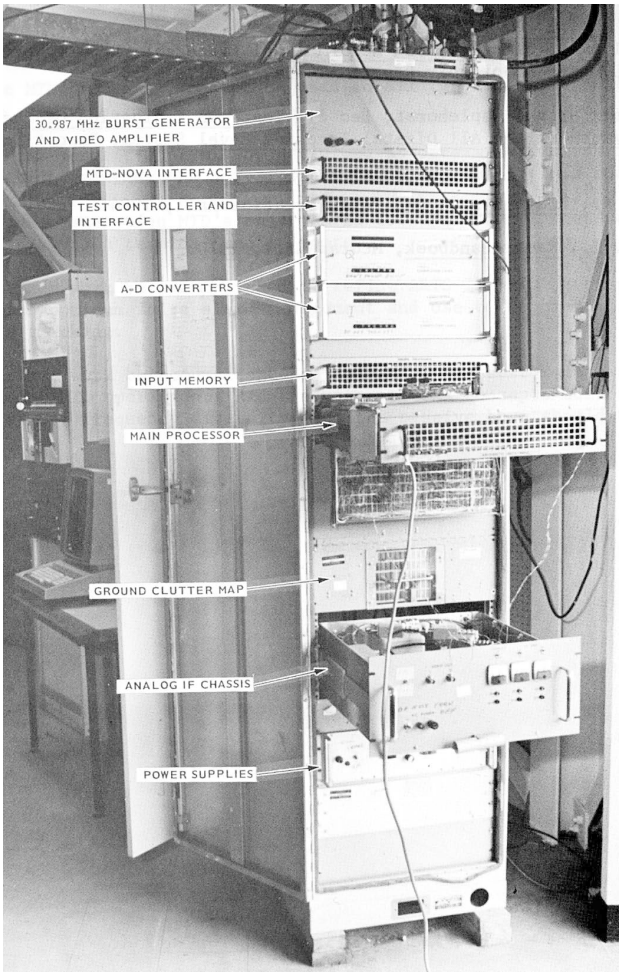

FIGURE 14. The moving-target detector (MTD-1) at the Federal Aviation Administration (FAA) facility in Atlantic City, New Jersey, in 1974. The MTD-1 was extensively tested in competition with a modern digitized version of the moving-target indicator (MTI) delay-line canceler.

amplitude and position and are used for target-track initiation and updating. Also in the post-processor, area thresholds are maintained to control excess false alarms, particularly from bird flocks. Each area of about sixteen square nautical miles is divided into several velocity regions. The threshold in each region is adjusted on each scan to achieve the desired limit on false alarms without raising the threshold so high that small aircraft are prevented from being placed in track status. The post-processor also implements a map of small areas, only a few resolution cells in extent, in which the clutter return is so high that false alarms occur repeatedly. Detection in these areas is censored.

The MTD-1 was initially tested at Lincoln Laboratory by using an S-band AN/FPS-18 radar with a klystron transmitter that had been modified to improve its stability. The MTD-1, which is shown in Figure 14, was transferred in late 1974 to the FAA's radar test facility near Atlantic City, New Jersey, where it was connected to an ARTS-3 radar. The FAA [57] and Lincoln Laboratory engineers tested the MTD-1 extensively. Figure 15 shows the results of radar detection tests of small aircraft in rain. Figure 15(a) shows the extent of rain clutter, and Figure 15(b) shows the detection and automatic tracking of a number of aircraft for about four minutes. The vertical track at the center is detection of automobiles on a road. Later improvements included automatic elimination of moving road vehicles. A competition [58] was held between the MTD and the RVD-4, which was an advanced version of the sliding-window detector that estimated the correlation of rain-clutter returns and readjusted the thresholds appropriately. In this competition the MTD radar's false-alarm and target-detection performances proved to be markedly superior to those of the RVD-4.

In December 1975, the U.S. Air Force Air Defense Command arranged to test the MTD-1 in the presence of active electronic countermeasures and chaff [59]. An Air Force EB-57 equipped with four hundred pounds of chaff along with swept, spot, and barrage jammers was used for the test. The EB-57 and another test aircraft were detected with nearly unity blip scan ratio as they flew through the chaff. These tests demonstrated the superior detection performance of the MTD-1 in chaff and jamming, accompanied by a low false-alarm rate.

With the establishment of the superior performance of these techniques in both military and civilian environments, it was not long before contractors were proposing using these techniques on most new air-defense radars and on new developments in air-traffic-control radars.

The MTD-2 and the Parallel Microprogrammable Processor

By 1975 the FAA had decided that the MTD class of radars was an effective solution to the problem of detecting aircraft in high-clutter environments, but

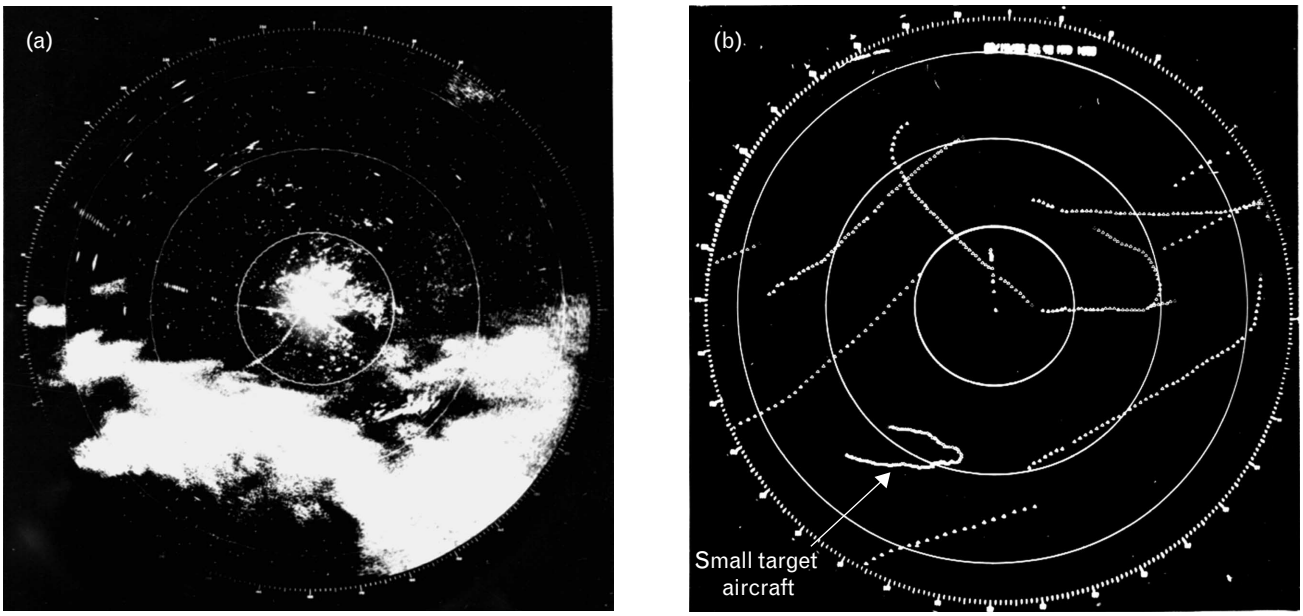

FIGURE 15. Performance of the MTD in heavy precipitation and ground clutter. This figure shows the detection of a small target aircraft in rain (a) with normal video, before the installation of the MTD, and (b) after the installation of the MTD. Notice the absence of false returns and the continuous tracking in the MTD image, even of aircraft with zero radial velocity. The target aircraft is a single-engine Piper Cherokee.

there were reservations concerning its complexity. Because the algorithms were embedded in the hardware, it would take a digital engineer or a highly trained radar technician to diagnose troubles. Lincoln Laboratory was encouraged to consider alternative designs that would relieve the logistic and maintenance problems that might arise. At that time, the concept of parallel processing was just evolving, and the notion that many signal processing problems lent themselves to architectures that applied a single, relatively rudimentary algorithm to multiple data sets was one of the innovative realizations of the power of digital signal processing. The parallel microprogrammed processor, or PMP [55, 60], was an important early example of this kind of architecture.

The PMP was an SIMD (single-instruction multiple-data-stream) computer consisting of a number of processing modules (typically two to eight), all served by one control unit. This type of system was seen as particularly appropriate for a surveillance radar such as the ASR, because the same algorithms are used for each range gate. One PMP module served ten nautical miles of range in an ASR. An extra processing module served as a spare, to be switched in when a fault was detected in the primary module.

A processing module consisted of two wire-wrapped boards: one to hold the input data and clutter-map memories; the other, the processing element, to handle all the mathematical computations. The processing element contained two 24-bit arithmetic and logic units, a bit shifter, and a small high-speed memory. The processing element operated with a 75-nsec instruction cycle, and on average it performed two simple operations per cycle time, resulting in a net processing rate of 25 million instructions per second. The control unit also consisted of two wire-wrapped boards. One board held memory for instructions, program constants, and target reports from the processing modules. Its processing element did all the required arithmetic, such as memory-address generation and time keeping, and interfaced with the processing modules and the post-processor. To handle this kind of computational workload, a PMP assembly language was developed at Lincoln Laboratory. Each line of code contained all the assembly language instructions to be executed in one cycle time. The machine language was generated by using a cross-compiler that was also written at Lincoln Laboratory

and executed on the Laboratory's central computer. Three PMP-1 devices were built at Lincoln Laboratory and seven PMP-2 devices, as shown in Figure 16, were built under contract by Stein Associates.

In 1978 Lincoln Laboratory installed an MTD-2, using a PMP-2 with the ASR-7 radar, at the Burlington, Vermont, airport [61]. The FAA chose this site near Mount Mansfield because it is reputed to have the worst clutter environment on the east coast. The FAA brought in air traffic controllers and other personnel from all over the United States to observe and operate the Laboratory's MTD-2 radar. Convinced that the MTD-2 was what they wanted, the FAA asked Lincoln Laboratory to help write the specifications for the next-generation ASR. A production contract was placed with Westinghouse for the ASR-9 radar, which today is in operation throughout the United States. The development of the MTD concept fundamentally changed the way surveillance radars are designed, and it caused that change quickly, essentially overnight! As a result of its successful implementation, the acronym MTD has since become an eponym.

FIGURE 16. The parallel microprogrammed processor (PMP) built by Stein Associates and installed at the Burlington, Vermont, airport in 1978. This detector was displayed to visiting air traffic controllers from all over the United States, who were positively impressed with its performance.

Space-Based Surveillance of the Earth

During the 1970s the Communications division at Lincoln Laboratory examined ways to reduce the vulnerability of military communication satellites to jamming. This need led to the study of adaptive-nulling techniques to minimize the effect of jamming. In 1985 the Laboratory began studying the feasibility of a large array radar that would search for moving ground or airborne targets from low-earth orbit. This proposed orbiting-radar design is another example of interdisciplinary technology transfer, showing how the expertise developed from the communications-satellite effort could be applied to problems in radar signal processing

A space-based surveillance radar must handle two major sources of interference: clutter from the entire visible earth and jamming in the antenna sidelobes. The clutter can be attenuated by using displaced-phase-center-antenna techniques (see the article by Muehe and Labitt in this issue); the sidelobe jammers can be attenuated by modifying some of the array-element weights to cause deep pattern nulls in the jamming directions. In both cases the formation of a single data stream from current and delayed data from many antenna elements requires the computation of a weighted sum for each sampling instant. This computation can require a very large number of multiplications and additions per second: four times the product of the number of antenna elements and the sampling rate. Because these operations are regular, it was determined that they could be implemented by using commercially available special-purpose integrated circuits.

The determination of the appropriate set of weights is another matter. These weights must be determined adaptively. As the satellite moves relative to the surface of the earth, each jammer appears to move

from one part of the sidelobe region to another, and therefore the weights must be readapted about two hundred times per second. An algorithm to compute these adapted weights is much more complicated than a simple sum of products, and in 1985 it seemed to require a computer capable of adding, subtracting, multiplying, dividing, computing square roots, and storing large amounts of data. At that time single-chip digital-signal-processing computers were available, but they were many times less efficient than the simple special-purpose chips for computing sums of products. The cost of carrying out the weight-adaptation algorithm depends sensitively on N, the number of weights being determined. The computational cost is proportional to the cube of N, so that determining the weights for twice as many antenna elements requires eight times the number of multiplications and additions.

Lincoln Laboratory engineers working on this problem in 1985 therefore estimated that it would be reasonable to fly enough computing power to adapt twenty-five weights, though there were many reasons why system designers might have wanted to use a larger number. In a very narrowband system with modest aperture, for example, $N + 1$ weights are required to null out N jammers. If the bandwidth of the radar is larger, or if the array aperture is large, several weights can be required per nulled jammer. Adaptation to clutter also requires many weights.

In the same year a small project was initiated to devise and demonstrate an efficient approach to the computation of adaptive weights. The result was the discovery, early in 1986, of an unique confluence of a technology, an algorithm, and an architecture that enabled the construction of an adaptive weighting computer called MUSE (Matrix Update Systolic Experiment). MUSE, a demonstration system, was capable of computing sixty-four weights several hundred times per second, but it had a physical size and weight no larger than a package of cigarettes. At that time, no actual adaptive antenna arrays with sixty-four elements existed: it would have made no sense to build such arrays, since nothing (i.e., no existing computer) could adapt their weights in real time.

The data used to determine the weights in the MUSE algorithm are a series of columns of complex numbers. Each column contains one sample from each of the N antenna elements. It is important to understand that the data arrive one datum at a time, one column at a time. This limited serial data transfer means that the number of data input pins required is quite reasonable.

The computation of the adaptive weights involves the triangularization of the raw data and a back-substitution that yields the actual weights. The triangularization process is in essence a sequence of two-dimensional rotations. These rotations are applied sequentially to the original data matrix until the matrix has all zeros in the upper-right portion and no zero values in the lower-left portion. The solution of the weights using back-substitution is then algorithmically straightforward and computationally simple.

Given that the critical part of the adaptive-weight computation can be reduced to a sequence of simple rotations, it became important to look for efficient ways to implement such a rotation. A design for such a rotating circuit was developed in the 1950s, and it is called a CORDIC module [47]. The CORDIC module is made up of adders and shifters, and it is easily pipelined so that it can accept new pairs of numbers as fast as it can add, even though any rotation takes much longer than any addition. A CORDIC module is a convenient size to be realized as a single integrated module. All ninety-six CORDIC modules required for MUSE are identical and can be easily interconnected. In this way the architecture of MUSE and the algorithm it carries out are perfectly adapted to each other.

A further improvement was the use of wafer-scale integration. This technology had been attempted by many laboratories in the 1980s, but Lincoln Laboratory was the first to succeed in building wafer-scale circuits [62]. The difficulty with wafer-scale integration is that even one tiny defect on a chip usually makes the chip nonfunctional. When the chip is a whole wafer, the probability of a defect becomes a virtual certainty. The Laboratory's approach was to build a wafer with redundant cells and to connect together enough of each type of cell to yield a working system. In the case of MUSE, there was only one type of cell, a CORDIC module. A wafer was fabricated with 132 CORDIC modules. Interconnections were made by

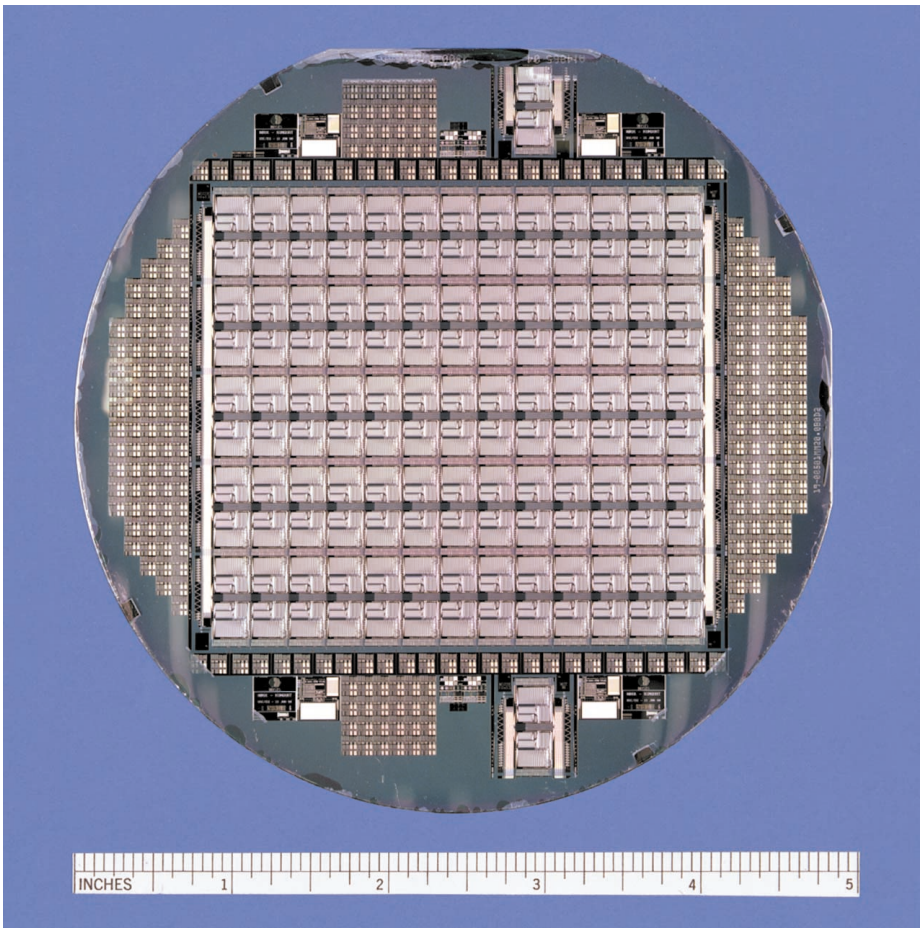

FIGURE 17. The MUSE (Matrix Update Systolic Experiment) wafer provided an efficient approach to the computation of adaptive weights. This demonstration system could compute sixty-four adaptive weights several hundred times per second.

using an automated laser weld to make electrical connections between the cells. The same automated laser was used to break connections, when necessary, by vaporizing metallization.

The active area of a MUSE wafer, shown in Figure 17, fit into a square of just over three inches on a side (or nine square inches in area). At a clock rate of 6 MHz, the system was able to carry out almost three hundred million rotations per second, equivalent to about three billion instructions per second in a conventional single-instruction computer. Power consumption was only about 10 W, and because there were so few wired connections, MUSE was a highly reliable design suitable for space applications. Through further refinement of the integrated-circuit fabrication technology, a modern version of MUSE developed by the Hughes Corporation has one thousand times the computational power of Lincoln Laboratory's original demonstration.

Summary

The proliferation of radar signal processing efforts at Lincoln Laboratory has been driven by the overwhelmingly dominant need to detect and measure fundamentally small radar-target returns in the presence of potentially overwhelming noise and other unwanted returns (i.e., clutter, both natural and intentional). This requirement has fundamentally involved the concurrent development of (1) theory and algorithms, (2) the underlying analog and digital technology [63], and (3) efficient architectures that merge theory and device technology into real systems for

important military and commercial applications—on the ground, in the air, and in space. These developments, which started with what might now be viewed as primitive efforts in SAGE and early ballistic missile defense, progressed through the development of fundamental device technology, both analog and digital, and have now moved in the direction of exploiting the enormous power and flexibility of digital processing, both custom and commercial.

For example, efforts are under way to develop extremely high-performance systems that combine classical clutter suppression with computationally challenging adaptive processing for joint detection of targets in clutter and jamming (a technique known as space-time adaptive processing, or STAP) [64]. Moreover, recent successes in radar imaging hold the promise for real-time and near-real-time generation of complex images that could be exploited by analysts for rapid adaptation to evolving circumstances. These combined techniques doubtlessly will find their way into future advanced ground, airborne, and space systems.

In viewing the history of signal processing, we note an interesting paragraph in Merrill Skolnik's 1962 seminal book on radar [65]: "The maximum compression ratios possible will depend upon the amount of development effort expended to achieve them. The numerical examples given by Krönert [66] for Gaussian-shaped pulses and cascaded-lattice networks indicate the feasibility of achieving pulse-compression ratios from $\beta\tau$ = 8 to 40. In Darlington's patent [67] an example is given for a Gaussian-shaped pulse in which a compression ratio of 34 is mentioned. The British patent issued to Sproule and Hughes [68] claims that it is possible to achieve a pulse-compression ratio of 100. Klauder [3] et al. also suggest that pulse-compression ratios of approximately 100 are possible." The extraordinary advances in radar signal processing in the past five decades admit technology that today allow radars with $\beta\tau$ significantly in excess of 1,000,000.

Acknowledgements

The efforts described in this article span fifty years of research at Lincoln Laboratory. As such it is impossible to acknowledge the efforts of all the people at the Laboratory, in government, and in industry who contributed to, supported, and encouraged these programs. The authors can only acknowledge the general support of the U.S. Army, the U.S. Air Force, U.S. Navy, DARPA (formerly ARPA), the Ballistic Missile Defense Organization (BMDO), and the FAA. However, James Carlson, now retired from the BMDO, deserves special mention for his long-term keen interest and broad support of radar signal processing, not only at Lincoln Laboratory, but across the country. Within the Laboratory, Irwin Lebow contributed enthusiasm and strength of leadership, and Ben Gold provided intellectual guidance in the field of digital signal processing at a time when it was just emerging. Lastly, the lead author would like to remember the late Jerry Margolin for his phenomenal insight and the incredibly spirited discussions that he provided a new young staff member at Lincoln Laboratory.

REFERENCES

1. Radiation Laboratory Series, L.N. Ridenour, ed., particularly vol. 24, *Threshold Signals,* J.L. Lawson and G.E. Uhlenbeck, eds. (McGraw-Hill, New York, 1950), pp. 28–29; CD-ROM (Artech House, Boston, 1999).
2. S.O. Rice, "Mathematical Analysis of Random Noise," pt. 1, *Bell Syst. Tech. J.* **23,** July 1944, pp. 282–332; pt. 2, *Bell System Tech. J.* **24,** Jan. 1945, pp. 46–160.
3. J.R. Klauder, A.C. Price, S. Darlington, and W.J. Albersheim, "The Theory and Design of Chirp Radars," *Bell Syst. Tech. J.* **39** (4), 1960, pp. 745–808.
4. A universal problem with signal processing in the early years, at Lincoln Laboratory and elsewhere, was the enormous lead that "theory" enjoyed over "practice." To quote Ed Kelly, "In the old days, we used to suggest various schemes based on theoretical work, but the hardware guys would just roll on the floor and hold their sides, because they were stuck with Rs, Ls, Cs, and quartz delay lines!"
5. T.C. Bartee, I.L. Lebow, and I.S. Reed, *Theory and Design of Digital Machines* (McGraw-Hill, New York, 1962).
6. G.P. Dinneen and I.S. Reed, "An Analysis of Signal Detection and Location by Digital Methods," *IRE Trans. Inf. Theory* **2** (1), 1956, pp. 29–38.
7. D.F. DeLong and E.M. Hofstetter, "On the Design of Optimum Radar Waveforms for Clutter Rejection," *IEEE Trans. Inf. Theory* **13** (3), 1967, pp 454–463.
8. D.F. DeLong and E.M. Hofstetter, "The Design of Clutter-Resistant Radar Waveforms with Limited Dynamic Range," *IEEE Trans. Inf. Theory* **15** (3) 1969, pp. 376–385.
9. B. Loesch, E.M. Hofstetter, and J.P. Perry, "A Technique for Synthesizing Signals and Their Matched Filters," *Lincoln Laboratory Technical Report 475* (29 Dec. 1969), DTIC #AD-704754.
10. E. Stern, "Microsound Components, Circuits, and Applications," *IEEE Trans. Microw. Theory Tech.* **17** (11), 1969, pp. 835–844.
11. E. Stern, "Microsound Components, Circuits, and Applications," *Ultrason.* 7 (4), 1969, pp. 227–233.
12. E. Stern and J.O. Taylor, "The Role of Acoustic Surface Wave Devices in Electronic Signal Processing," *1971 IEEE-GMTT Int. Microwave Symp. Dig., Washington, 16–19 May 1971*, pp. 48–50.
13. H.M. Gerard, W.R. Smith, W.R. Jones, and J.B. Harrington, "The Design and Applications of Highly Dispersive Acoustic Surface-Wave Filters," *IEEE Trans. Sonics & Ultrason*. **20** (2), 1973, pp. 94–104.
14. J.D. Maines and E.G.S. Paige, "Surface-Acoustic-Wave Devices for Signal Processing Applications," *Proc. IEEE* **64** (5), 1976, pp. 639–652.
15. T.A. Martin, "The IMCON Pulse Compression Filter and Its Applications," *IEEE Trans. Sonics Ultrason.* **20** (2), 1973, pp. 104–112.
16. The "linear frequency modulated" (LFM) waveform has been the workhorse of radar signals since its development in World War II. It allows a peak-power-limited transmitter to improve detectability (energy) by expanding the pulse length and, while sweeping frequency and thereby increasing the bandwidth, retain the range-resolution properties of a short pulse. A range-Doppler ambiguity is introduced, but this is often either not an issue or surmountable by separate measurements and/or *a priori* knowledge.
17. C.E. Cook and M. Bernfield, *Radar Signals*: *An Introduction to Theory and Application* (Academic Press, New York, 1967).
18. In the early 1960s, Jerry Freedman asked Jerry McCue to look at the signals that bats ("little brown and big brown") use to locate targets. It appears that bats also use frequency modulation, in their case roughly hyperbolic, to achieve both detectability and range resolution. An interesting paper, written from a radar signal processing perspective, is by J.J.G. McCue, "Aural Pulse Compression by Bats and Humans," *J. Acoust. Soc. Am.* **40** (3), 1966, pp. 545–548.
19. R.C. Williamson and H.I. Smith, "The Use of Surface-Elastic-Wave Reflection Gratings in Large Time-Bandwidth Pulse-Compression Filters," *IEEE Trans. Sonics & Ultrason.* **20** (2), 1973, pp. 113–123.
20. M. Axelbank, W.W. Camp, V.L. Lynn, and J. Margolin, "ALCOR—A High-Sensitivity Radar with One-Half-Meter Range Resolution," *IEEE 1971 Int. Conv., New York, 22–25 Mar. 1971,* pp. 112–113.
21. W.J. Caputi, Jr., "Stretch: A Time-Transformation Technique," *IEEE Trans. Aerosp. Electron. Syst.* 7 (4), 1971, pp. 268–278.
22. Also applied to a stepped FM technique in TRADEX. Private communication.
23. R.C. Williamson, "Properties and Applications of Reflective-Array Devices," *Proc. IEEE* **64** (5), 1976, pp. 702–710.
24. V.S. Dolat and R.C. Williamson, "BGO Reflective Array Compressor (RAC) with 125 μs of Dispersion," *IEEE Ultrasonics Symp. Proc., Los Angeles, 22–24 Sept. 1975,* pp. 390–394.
25. J.M. Smith, E. Stern, and A. Bers, "Accumulating-Layer Surface-Wave Convolver," *Electron. Lett.* **9** (6), 1973, pp. 145–146.
26. J.H. Cafarella, J.A. Alusow, N.M. Brown, and E. Stern, "Programmable Matched Filtering with Acoustoelectric Convolvers in Spread-Spectrum Systems," *Ultrasonics Symp. Proc., Los Angeles, 22–24 Sept. 1975*, pp. 205–208.
27. K.A. Ingebrigtsen and E. Stern, "Coherent Integration and Correlation in a Modified Acoustoelectric Memory Correlator," *Appl. Phys. Lett.* **27** (4), 1975, pp. 170–172.
28. R.W. Ralston, J.H. Cafarella, S.A. Reible, and E. Stern, "Improved Acoustoelectric Schottky Diode/$LiNbO_3$ Memory Correlator," *Ultrasonics Symp. Proc., Phoenix, Ariz., 26–28 Oct. 1977,* pp. 472–477.
29. A fascinating footnote to history reveals that Hedy Lamarr, the famous actress, invented a form of frequency hopping during World War II. H.-J. Braun, "Advanced Weaponry of the Stars," *Am. Heritage Mag. of Invention and Technol.* **12** (4), 1997, pp. 10–16; R. Price, "Further Notes and Anecdotes on Spread-Spectrum Optics," *IEEE Trans. Commun.* **31** (1), 1983, pp. 85–97; D.R. Hughes and D. Hendricks, "Spread-Spectrum Radio," *Sci. Am.* **278** (4), 1998, pp. 94–96.
30. R.C. Williamson, V.S. Dolat, R.R. Rhodes, and D.M. Boroson, "Satellite-Borne Chirp-Transform System for Uplink Demodulation of FSK Communication Signals," *IEEE Ultrasonics Symp. Proc., New Orleans, 26–28 Sept. 1978,* pp. 741–747.
31. D.T. Bell and R.C.M. Li, "Surface-Acoustic-Wave Resonators," *Proc. IEEE* **64** (5), 1976, pp. 711–721.
32. D.L. Spears and H.I. Smith, "High-Resolution Pattern Replication Using Soft X-Rays," *Electron. Lett.* **8** (4), 1972, pp. 102–104.
33. C.M. Rader and B. Gold, "Digital Filter Design Techniques in the Frequency Domain," *Proc. IEEE* **55** (2), 1967, pp. 149–171.
34. B. Gold and C. Rader, *Digital Processing of Signals* (McGraw-Hill, New York, 1969).

35. R.J. Purdy, P.E. Blankenship, A.E. Filip, J.M. Frankovich, A.H. Huntoon, J.H. McClellan, J.L. Mitchell, and V.J. Sferrino, "Digital Signal Processor Designs for Radar Applications," *Technical Note 1974-58,* Lincoln Laboratory (31 Dec. 1974).
36. The digitization process is crucial to the implementation of the processing. At the time, analog-to-digital (A/D) converters did not exist with the requisite word length (dynamic range) and sample rate. Consequently, Lincoln Laboratory initiated a development with Hughes Aircraft for a 10-bit, 60-Msec/sec A/D converter. This effort was highly successful. Four A/D pairs (I and Q) were built. Eventually, one pair was used on the Army's Signature Measurements Radar and two were used for many years on the Cobra Judy Radar. (See the article entitled "Radars for Ballistic Missile Defense Research," by Philip A. Ingwersen and William Z. Lemnios, in this issue.)
37. B. Gold, I.L. Lebow, P.G. McHugh, and C.M. Rader, "The FDP, a Fast Programmable Signal Processor," *IEEE Trans. Comput.* **20** (1), 1971, pp. 33–38.
38. L. Rabiner and B. Gold, *Theory and Application of Digital Signal Processing* (Prentice-Hall, Englewood Cliffs, N.J., 1975).
39. E.M. Hofstetter, personal communication, Lincoln Laboratory, ca. April 1974.
40. J.A. Feldman, E.M. Hofstetter, and M.L. Malpass, "A Compact, Flexible LPC Vocoder Based on a Commercial Signal Processor Microcomputer," *IEEE J. Solid-State Circuits* **18** (1), 1983, pp. 4–9.
41. C.E. Muehe, Jr., L. Cartledge, W.H. Drury, E.M. Hofstetter, M. Labitt, P.B. McCorison, and V.J. Sferrino, "New Techniques Applied to Air-Traffic Control Radars," *Proc. IEEE* **62** (6), 1974, pp. 716–723.
42. B. Gold and C.E. Muehe, "Digital Signal Processing for Range-Gated Pulse Doppler Radars," *XIXth AGARD Conf. Proc. on Advanced Radar Systems, No. 66, 25–29 May 1970, Istanbul, Turkey.*
43. One key advantage of digital processing is the ability to precisely simulate the computations in advance. This approach was used extensively in the design of the Digital Convolver System.
44. A.H. Anderson, J.M. Frankovich, L. Henshaw, R.J. Purdy, and O.C. Wheeler, "The Digital Convolver System," *Project Report SDP-228,* Lincoln Laboratory (19 June 1981).
45. P.E. Blankenship and E.M. Hofstetter, "Digital Pulse Compression via Fast Convolution," *IEEE Trans. Acoust. Speech Signal Process.* **23** (2) 1975, pp. 189–201.
46. B. Gold and T. Bially, "Parallelism in Fast Fourier Transform Hardware," *IEEE Trans. Antennas Propag.* **21** (1), 1973, pp. 5–16.
47. J.E. Volder, "The CORDIC Trigonometric Computing Technique," *IRE Trans. Electron. Comput.* **8** (3), pp. 330–334.
48. As part of an Air Force–supported internal Lincoln Laboratory effort, the Laboratory developed and fabricated a set of very fast 2-bit adder/subtractor circuits that was ideally suited for fast signed arithmetic array multiplication. This design was modified by Peter E. Blankenship into a single programmable adder/subtractor component suitable for FFT and CORDIC rotator computations. This design was subsequently transferred by the Laboratory to Motorola and incorporated in their MECL10K product line as the MC10287L.
49. S.D. Pezaris, "A 40-ns 17-Bit by 17-Bit Array Multiplier," *IEEE Trans. Comput.* **20** (4), 1971, pp. 442–447.
50. H.L. Groginski and G.A. Works, "A Pipeline Fast Fourier Transform," *EASCON Record, Washington, 27–29 Oct. 1969*, pp. 22–29.
51. G.C. O'Leary, "Nonrecursive Digital Filtering Using Cascade Fast Fourier Transforms," *IEEE Trans. Audio Electroacoust.* **18** (2), 1970, pp. 177–183.
52. The DCS was designed by Lincoln Laboratory and manufactured by General Electric Heavy Military Systems Group to a detailed design specification.
53. W.H. Drury, "Improved MTI Radar Signal Processor," *Lincoln Laboratory Project Report ATC-39* (3 Apr. 1975), FAA-RD-74-185, DTIC #ADA-010478/6.
54. L. Cartledge and R.M. O'Donnell, "Description and Performance Evaluation of the Moving Target Detector," *Lincoln Laboratory Project Report ATC 69* (3 Aug. 1977), DTIC #ADA-040055.
55. W.H. Drury, B.G. Laird, C.E. Muehe, and P.G. McHugh, "The Parallel Microprogrammed Processor (PMP)," *Radar '77, London, 25–28 Oct. 1977.*
56. W.W. Schrader and V. Gregers-Hansen, "MTI Radar," chap. 15, *Radar Handbook,* M.I. Skolnik, ed., 2nd edition (McGraw Hill, New York, 1990), pp. 15.1–15.72.
57. R.S. Bassford, W. Goodchild, and A. DeLaMarche, "Test and Evaluation of the Moving Target Detector," Final Report, Oct. 1977, FAA-RD-77-118, DTIC #ADA-047887.
58. R.M. O'Donnell and L. Cartledge, "Comparison of the Performance of the Moving Target Detector and the Radar Video Digitizer," *Project Report ATC-70,* Lincoln Laboratory (26 Apr. 1977), NTIS No. ADA-040472.
59. R.M. O'Donnell and L. Cartledge, "Evaluation of the Performance of the Moving Target Detector (MTD) in ECM and Chaff," *Technical Note 1976-17,* Lincoln Laboratory (25 Mar. 1976).
60. G.P. Dinneen and F.C. Frick, "Electronics and National Defense: A Case Study," *Science* **195** (4283), 1977, pp. 1151–1155.
61. D. Karp and J.R. Anderson, "Moving Target Detector (Mod II) Summary Report," *Project Report ATC-95,* Lincoln Laboratory (3 Nov. 1981), DTIC #ADA-114709.
62. C.M. Rader, "Wafer-Scale Integration of a Large Systolic Array for Adaptive Nulling," *Linc. Lab. J.* **4** (1), 1991, pp. 3–30.
63. An interesting snapshot of the state of the art of both analog and digital processing technology, circa 1977, is contained in chap. 10, by R.J. Purdy, of *Radar Technology,* E. Brookner, ed. (Artech House, Dedham, Mass., 1977), pp. 155–162.
64. J. Ward, "Space-Time Adaptive Processing for Airborne Radar," *Lincoln Laboratory Technical Report 1015,* Dec. 1994, DTIC #ADA-293032.
65. M.I. Skolnik, *Introduction to Radar Systems* (McGraw-Hill, New York, 1962), p. 495.
66. R. Krönert, "Impulsverdicktung [Pulse Compression]," pt. 1, *Nachr. Tech. Elektron.* **7**, Apr. 1957, pp. 148–152, 162; pt. 2, *Nachr. Tech. Elektron.* **7**, July 1957, pp. 305–308. For English abstractions of these two articles, see abstract 72, *Proc. IRE* **46** (2),1958; abstract 1078, *Proc. IRE* **46** (5), 1958, p. 936.
67. S. Darlington, "Pulse Transmission," U.S. Patent No. 2,678,997, 18 May 1954.
68. "Improvements in and Relating to Systems Operating by Means of Wave Trains," British Patent Specification 604,429, 5 July 1948, issued to Henry Hughes and Sons, Ltd., D.O. Sproule, and A.J. Hughes.

ROBERT J. PURDY received an S.B. degree from MIT, and M.S. and Ph.D. degrees from Purdue University, all in electrical engineering. In 1968 he joined the Radar Systems group at Lincoln Laboratory, and in 1972 he was appointed assistant leader of the Digital Signal Processing group. He was subsequently appointed associate leader and then leader of the newly formed Sensor Processing Technology group. In this capacity, he directed the design and fabrication of several state-of-the-art radar signal processors. He also helped identify the need for custom integrated circuits with specific application to radar signal processing, and several of these components were later fabricated by the Integrated Circuit group. He has published several articles on radar signal processing, notably chapter 10 of *Radar Technology*, edited by E. Brookner, and chapter 5, with J.H. McClellan, of *Applications of Digital Signal Processing Technology*, edited by A.V. Oppenheim. He also published, with G. Heiligman, an article on an adaptive beamforming processor. For three years he served as Radar Editor for the *IEEE Transactions on Aerospace and Electronic Systems*.

PETER E. BLANKENSHIP received S.B.E.E., S.M.E.E., and Electrical Engineer degrees from MIT. He joined Lincoln Laboratory in 1968 as a technical staff member and was appointed assistant leader of the Speech Systems Technology group in 1979. He was appointed associate leader in 1981, and became associate head of the Computer Technology division in 1984. In 1995 he joined the Surveillance and Control division as an associate division head, with responsibilities in the Information System Technology area. In 1998 he was appointed senior staff in the Communications and Information Technology division office. Later that year he was assigned as senior staff in the Director's Office, with responsibilities in the areas of technology transfer, intellectual property, continuing technical education, and information technology infrastructure. His research interests include high-performance computing architectures, digital-signal-processing algorithms and systems, speech processing systems, efficient digital system implementations, novel applications of VLSI and wafer-scale integration circuit technology, and medical electronics. He is a member of Tau Beta Pi and Eta Kappa Nu, and a Senior Member of the IEEE. He received the IEEE Centennial Medal in 1984 for service to the Signal Processing Society, where he served on the Conference Board for nine years.

CHARLES EDWARD MUEHE received a B.S. degree from Seattle University in 1950 and an S.M. degree from MIT in 1952, both in electrical engineering. After teaching at Seattle University for four years, he joined the Microwave Components group at Lincoln Laboratory. Of the microwave systems he helped develop the most notable were used on the Laboratory's Haystack planetary radar in the fourth test of Einstein's general theory of relativity, and on the ALCOR radar at the Kwajalein Atoll missile test range, which is capable of imaging reentry vehicles. In 1967 he became associate leader, and in 1968 group leader, of a group that designed, built, and tested complete prototype radar systems. The first system was a radar used in Vietnam to detect people walking under dense foliage. Starting in 1972 his group developed digital signal and data processors capable of completely automatic detection, tracking, and displaying of moving targets in heavy clutter. This work led to a netted radar system demonstrated by the Army Artillery at Fort Sill, Oklahoma, an airborne radar to detect slowly moving ground vehicles (the progenitor of the Joint STARS radar), and the prototype of a widely employed FAA Airport Surveillance Radar (ASR-9). For seven years he served as Radar Editor of the *IEEE Transactions on Aerospace and Electronic Systems*.

CHARLES M. RADER
is a senior staff member in the Embedded Digital Systems group. He received a B.E.E. degree and an M.E.E. degree in electrical engineering from the Polytechnic Institute of Brooklyn. His many accomplishments include research on speech bandwidth compression, contributions to the field of digital signal processing, application of optical techniques to educational technology and communication, and investigation of space-based radar systems. He has been at Lincoln Laboratory since 1961. From 1971 to 1982 he was an assistant group leader in the Spacecraft Processors group, which built the LES-8 and LES-9 communications satellites launched in March 1976. He is a Fellow of the IEEE and past president of the IEEE Acoustics, Speech and Signal Processing (ASSP) Society. He has received the ASSP Technical Achievement Award (1976), the ASSP Society Award (1985), and the IEEE Jack S. Kilby Award (1996). He has also taught courses on advanced digital signal processing in China and Mexico. His books include *Digital Processing of Signals* (coauthored with Ben Gold), *Number Theory in Digital Signal Processing* (coauthored with James McClellan), *Advanced Digital Signal Processing* (coauthored with John G. Proakis, Chrysostomos L. Nikias, and Fuyun Ling) and *Digital Signal Processing* (coedited with Lawrence Rabiner).

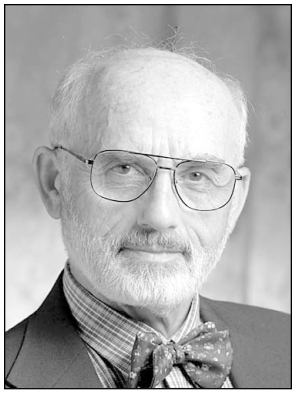

ERNEST STERN
was associate head of the Solid State division from 1983 to 1995. As leader of the Analog Device Technology group from 1969 to 1983, he led and participated in the development of surface-acoustic-wave (SAW) and superconductive-device technology as well as X-ray lithography and precision ion-beam-etching equipment. The SAW activity produced unique, precisely matched pulse compressors for the ALCOR all-range channel and for the multiple-antenna surveillance radar (MASR), as well as enabling devices for military satellite (MILSAT) communication prototypes. Additional devices were developed and incorporated into prototypes of high-performance spread-spectrum IFF systems, a packet-radio system, and a missile-borne antijamming system. Before joining Lincoln Laboratory in 1964, he developed ferrite microwave devices at Sperry Gyroscope and General Electric Companies, and was director of research at Microwave Chemical Laboratory. He served in the U.S. Navy from 1946 to 1948, received a B.S. degree in electrical engineering from Columbia University in 1953, and attended the Graduate School of Electrical Engineering at Cornell University from 1953 to 1955. He retired from the Laboratory in March 1995.

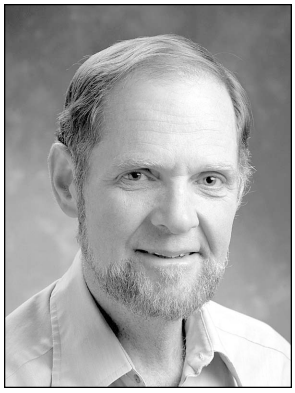

RICHARD C. WILLIAMSON
is a senior staff member in the Electro-Optical Materials and Devices group, where he carries out research and development of a wide range of optical devices and subsystems. He received S.B. and Ph.D. degrees in physics from MIT. His Ph.D. thesis work on ultrasonic studies of low-temperature physics was followed by a five-year postdoctoral position in the MIT Materials Science Center. He joined Lincoln Laboratory in 1970 and worked on the development of SAW devices for signal processing in radar and communication systems. He led the development of the first reflective-array compressors and became associate leader of the Analog Device Technology group. His transition from basic research to systems development included a trip to the Kwajalein Atoll radar-measurements field site, where he debugged signal processing hardware and installed wideband SAW devices in the ALCOR radar. In 1980, he became the leader of the Applied Physics group, which carried out development of optical devices. He has published three book chapters and many technical papers on a wide range of topics. Professional activities include considerable involvement with the Optical Society of America and the IEEE. He is an IEEE Fellow, received a Career Achievement Award, and was a National Lecturer.