Wide-Bandwidth High-Resolution Search for Extraterrestrial Intelligence

Semiannual Status Report

15 June 1992 -- 15 Dec 1992

Paul Horowitz, Principal Investigator

Harvard University
Cambridge, MA
02138

June 15, 1992

Grant Number NAGW-2872

1. INTRODUCTION

This report summarizes research accomplished during the second 6-month period of the grant. During the period covered by this report the active personnel included the PI, one graduate student (Darren Leigh), two Harvard undergraduates (Greg Galperin and Derrick Bass), and a recent mathematics graduate from Harvard (Nick Shectman).

2. RESEARCH ACCOMPLISHED

2.1 Antenna Configuration

Our new SETI system is designed to reject terrestrial interference (the #1 problem in SETI) by exploiting the property that a genuine extraterrestrial transmitter must be both pointlike and exhibit sidereal rotation. Thus we are building a two-horn receiver, with stationary beam lobes oriented east-west, along with a third "terrestrial" low-gain antenna. Although we originally envisioned a pair of beams separated by several beam widths, we now favor some degree of overlap (suggested by Prof. Staelin at MIT), such that the handoff from one beamlobe to the other keeps the source in sight continuously. To implement this, we looked at two schemes (Figure 1), namely *i)* a phased array consisting of 10 hexagonally-packed feedhorns and three low-noise preamps (with the central cluster of 4 horns passively combined, then buffered and phased with each of the passively combined outer sets of three horns; each horn could be either linear or dual-circular polarization), and *ii)* a simpler arrangement of two pyramidal (linear polarization) horns, aligned along their E-plane axes. The 10-horn arrangement exhibits a remarkable azimuthal symmetry of beam pattern (Figure 2), but requires a major (lengthy and expensive) construction effort; by contrast, the pyramidal pair is sensitive to only one linear polarization, but is easy to build and try out.

With either scheme we were told that near-field aperture interactions would cause major distortions of the far-field pattern, though no one was able to quantify the effect. To satisfy our curiosity we made some laboratory "test-range" measurements with a pair of pyramidal x-band (3 cm) horns, driven (via a magic-T hybrid) both alternately and simultaneously with Gunn oscillator sources, while scanning the far-field pattern with a small dipole connected to a spectrum analyzer. We were unable to see any interaction effects at the measurement accuracy (approximately 1dB).

While favoring the simplicity of the stacked pyramidal horns, we were concerned about two additional issues: *i)* does the far-field pattern have reasonable azimuthal symmetry, and *ii)* is there adequate beam overlap with separate horns (which is guaranteed with the interleaved 7-horn array)? To answer these questions we performed diffraction calculations, starting with a pyramidal horn design whose H-plane to E-plane dimensions are in the ratio of 1.35 (this ratio produces equal -3dB beam widths, owing to cosine taper in the H-plane field amplitude, combined with uniform field amplitude along the E-plane). We assumed no feedhorn interaction, and simply calculated the far-field pattern, using the parameters of our 26-meter antenna (full-width illumination angle of 18.5 degrees). Figure 3 shows the far-field beam intensity, from a single displaced horn, for three choices of horn aperture. In each case the horn center has been offset along the E-plane, relative to the Cassegrain axis, by half the horn aperture

(i.e., stacked horns). A rule of thumb to achieve maximum efficiency in Cassegrain design is to taper the illumination to approximately -10dB at the reflector edge. That corresponds to Figure 2c, producing beam overlap at the -6dB point; it also results in feedhorns that do not fit in the radome!

We finally settled on the design of Figure 2b, which we estimated to have a paraboloid efficiency (i.e., spillover efficiency times aperture efficiency) just 0.6dB less than the ideal; its taper at the edge of the dish is -5.5dB, compared with the conventional -10dB, resulting in somewhat increased sidelobe amplitude. It will fit in the radome, if the outside corners are cut diagonally (Figure 4). The feedhorn length was chosen to produce wavefront curvature of about 0.2 by 0.3 wavelengths (E-plane by H-plane), resulting in finished feedhorns that fit in the radome with about 1 inch to spare (Figure 5). They are constructed of 1/8" aluminum sheet (6061-T6), heliarc welded and joined to a WR-650 waveguide section with flange.

We mounted the horns, and performed drift scans of astronomical points sources (Sgr A, Cyg A). Figure 6 shows such a scan. The beam shape and overlap are ideal. However, the observed signal strength is lower than expected. With the help of absorber material (kindly provided by John Kraus), we expect to track down the problem when warm weather returns.

2.2 HEMT Low-Noise Amplifiers

We completed the first of the wideband (1.3 GHz to 1.9 GHz) low-noise L-band HEMT (high electron mobility transistor) amplifiers, using a kit of parts from Berkshire Technologies (Oakland, CA). Figure 7 shows the completed amplifier (cover removed), and Figure 8 shows its performance. It achieves 7K noise temperature over the waterhole band (1.4-1.7 GHz) when cooled to liquid nitrogen temperature, and 30K when operated at room temperature. We expect to complete the three remaining amplifiers (3 beams total, plus one spare) within the current 6-month period of the grant. Given the amplifier's excellent room-temperature noise performance (a factor of 2 better than our currently operating GaAsFET amplifiers), we intend to postpone the engineering of a cryogenic dewar until we have had some experience with the full operating system, which will be carried out with uncooled amplifiers. For ultimate cryogenic operation we favor the compact and convenient dewars from Infrared Laboratories (Tucson, AZ).

The use of amplifier kits (rather than completed amplifiers) results in a cost savings of $20k.

2.3 Downconverter

The downconverter subassemblies are nearly complete. The channelizing downconverter uses an array of 20 local oscillators (LO's), spanning the 40-80 MHz range in steps of 2 MHz, and phase-locked to the 10 MHz GPS-disclipined master station oscillator. The individual LOs feed a mixer-filter-digitizer board array, each of which accepts three IF inputs (east horn, west horn, and terrestrial antenna), and produces corresponding complex digitized baseband signals to feed the FFT array.

2.3.1 Local Oscillator Array

As described in the previous progress report, we explored several oscillator approaches (DDS, crystals, filtered comb, phase-locked loop), and wound up with a discrete varactor-tuned JFET oscillator phase-locked loop synthesizer, controlled by the elegant MC145170. We put ten oscillators on each of two boards, downloaded by a 87C751 microcontroller, and locked to a common 10 MHz station clock. Performance is excellent, with extremely low phase noise and spurs, as documented in the previous report.

We have now completed and tested the LO array (Figure 9), which works flawlessly. To give this synthesizer a front panel worth looking at, we built an array of 20 LED bar graph displays (also in Figure 9; schematic in Figure 10), showing the VCO loop voltage of the individual oscillators.

2.3.2 Mixer-Filter-Digitizer

The LO array drives the mixer-filter-digitizer subsystem, which is now completely designed and ready for board fabrication. The design has been modified somewhat since the previous report, because our FFT simulations demonstrate that our novel "dc autozero" circuit, though elegant from a circuit design standpoint, is unnecessary!

The final design (Figure 11) performs well, and in addition economizes in several ways: We changed our design to use inexpensive component amplifiers (Mini-Circuits Labs MAV-11) for the LO buffer, instead of the industry standard Avantek GPD amplifiers. We also decided to use discrete 7-pole anti-alias filters (less than $5 each), rather than the cheapest commercial units (TTE Inc., about $15). Another interesting feature of this design is the use of current-feedback baseband amplifiers (Analog Devices AD846), which are extremely stable and low-noise (2 nV/$\sqrt{\text{Hz}}$, and which have a compensation pin that can be used for dc clamping. We used that feature to clamp the ADC driver, thus permitting dc coupling to a CMOS converter (TMC1175) without fear of SCR latchup. Note the use of a floating precision 2.5V reference (the AD680), setting the conversion range to 0.6V-2.8V, which keeps the Schottky-clamped input between the supply rails; dc offset trim, referenced to this voltage range, is now performed at the driver's input, rather than at the ADC. The digitized output is three-stated onto a common 8-bit bus as complex pairs.

Figure 12 shows the completed printed circuit design, an economical 2-sided ground-plane board accepting the three IF inputs and single LO drive. We expect to fabricate and test 20 of these boards during the current 6-month grant period.

2.4 FFT Array

As described in the previous report, we have devised a 3-chip design for a 4 megapoint complex FFT, using the Austek A41102. The long transform is implemented as a succession of shorter row and column FFTs, with complex ("twiddle") multiplies interposed between the shorter transforms during the "corner turns." The Austek chip, originally designed for radioastronomy applications, is ideal for such continuous-flow real-time transforms; in addition, it requires no

special memory, and has great flexibility in terms of bit scaling, word width, transform length, normal vs bit-reversed sequences, and the use of an internal multiplier and data switches.

2.4.1 Simulations

We devoted considerable time to simulating our FFT architecture, using a simulator ("FDPSIM") supplied by Austek. In particular, *i)* we needed to verify that the 3-chip architecture is numerically correct, *ii)* we wished to determine the effect of word size quantization and roundoff on the spectral dynamic range, *iii)* we wished to determine the effects of finite twiddle factor word size ("depth") and quantization ("width"), since the use of full-size twiddle factor ROMs (e.g., 20 bits by 1M points) would raise costs considerably, and *iv)* we wished to verify by actual numerical simulation that a weak sinusoidal signal embedded in wideband noise, in the presence of additional strong sinusoidal signals, could be reliably and accurately detected by the FFT system we intend to build. We now describe the results of these simulations in detail, because there appears to be considerable confusion (and not a little folklore and mythology) in the signal processing community on precisely this issue. We hope the reader will be as surprised and enlightened by these results as we ourselves were.

Austek's simulator was never intended for such large transforms, and, as supplied, it took nearly a day to complete a single 4 megapoint transform on a Sun SPARC-2 (it took many days on a '486-type PC). We began by verifying the numerical correctness of the 3-chip architecture on scaled-down transforms, then proceeded to modify some of the simulator's modules to speed up performance; 4 million point simulations now take 2 hours.

The first simulations verified that a suite of sinusoidal waves, chosen with frequencies relatively prime but with each sine "on-bin" (i.e., an integer multiple of the lowest FFT frequency, $1/T_{fft}$), and covering a range of amplitudes, was properly resolved by the FFT, when using "perfect" (double precision floating point) arithmetic and twiddle factors.

We then explored the effects of finite word length in the FFT computation, in particular the 16-bit and 20-bit integer options that can be set by the initial command register load of the A41102. The results can be summarized as follows: With all "scales" enabled (i.e., with a 1-bit right shift of data following each FFT butterfly, required to prevent word growth for coherent frequencies present in the initial time series) the effect of finite word length and arithmetic precision is to introduce a "numeric noise" into the spectrum, consisting of an average of 1 LSB fluctuations in the final spectral amplitudes (and a *peak* fluctuation of 1.4 LSB, i.e., 1 LSB in each of the real and imaginary components). Stated this way, the result is independent of word length. It may seem surprising that a 2^{22}-point integer computation of the Fourier Transform introduces so little roundoff error; but the effect of the successive scale-by-2's is to keep pushing the roundoff error off the right end of the word.

The numeric noise is, of course, to be compared with any periodic signal present in the digitized input. A single tone, present as a full-scale on-bin sinusoid in the digitized time series, produces a full-scale output. Thus a 16-bit "all-scale" integer transform has a dynamic range of 2^{15} in amplitude (90dB); as we will soon see, however, there are several effects that can introduce spurious responses in the spectrum from a single sinusoidal input. These "spurs" can be

important in SETI, because a single interfering signal may produce a set of false responses in addition to the obvious large peak. Some of these effects are finite input quantization (i.e., shorter than the computation word size), spectral "leakage" (signal "off-bin"), and truncated twiddle ROM (both in word size and argument step size). We discuss these in the following paragraphs.

The effect of omitting some of the right-shift scales is interesting: Most obviously, one introduces a risk of numeric overflow -- a full-scale input sine causes overflow if any scales are omitted, a half-scale input causes overflow if more than one scale is omitted, etc. I.e., the spectral amplitudes grow by a factor of 2 for each omitted scale. This seems obvious, and in fact one might easily conclude that roundoff error grows the same way. However, the situation is more complicated -- it turns out, as revealed by our simulations, that the *peak* numeric noise grows as expected (a factor of 2, or one bit, in amplitude for each missing scale), but the *average* numeric noise amplitude grows only as the square root of the number of omitted scales (1/2 bit per missing scale). Thus one can squeeze some extra average dynamic range out of an integer FFT by omitting some scales, at the risk of numeric overflow (if a large signal is present); but note that the dynamic range relative to the *peak* numeric noise is not improved.

We studied the effects of ROM truncation next. To set the stage, note that a "full-sized" complex twiddle ROM (4M x 20 bits, say) would require 40 4-megabit ROMs, currently priced at about $30 each, thus approximately doubling the parts cost! Of course, one need not store both sine and cosine (factor of 2 savings), and one need store only a quarter-sine table (another factor of 4); that puts the ROM cost at about $150. Even at that price the ROMs are a significant portion of the board cost, so it is worth asking how wide and deep the ROM needs to be.

We ran a set of simulations, and learned the following: *i)* the spectral amplitude of an on-bin sinusoidal signal is very little affected by rather extreme ROM wordsize truncation; in particular, 8-bit ROM amplitudes affect spectral amplitudes by less than 1%; *ii)* ROM "width" (number of table entries) can also be reduced substantially, with almost negligible effect upon signal amplitude, but with production of spurs that are absent when using a full-sized ROM; *iii)* if ROM width is to be reduced for a non-square corner turn, truncate the larger address first.

Further explanation of *ii)* and *iii)*: Our 4M-point transform is implemented as 128x128x256, with a "small" (16K) twiddle factor multiplication following the first corner turn, and a "large" (4M) twiddle factor multiplication following the second corner turn (there are, in addition, a pair of 4M corner turns, without twiddle multiplication, at both ends of the overall FFT). The small twiddle factor ROM is cheap, and no truncation is needed there. The second ROM is the issue. It is a 16Kx256 corner turn (the initial 128x128 transform pair is exactly equivalent to a single 16K transform), requiring 14 and 8 address bits, respectively (22 address bits for a full-size ROM: 4M coefficients). We found that one can use a 16-bit (amplitude) ROM with 8 bits by 8 bits of address (256x256, or 64K complex coefficients, a factor of 64 less than a full-sized ROM) with no loss of signal amplitude, but with production of spurs whose peak amplitude is -54dBc (dB relative to the "carrier," i.e., the sinusoidal signal). That peak spur occurs at $f_o \pm 16K$ bins, with additional spurs at multiples of 16K bins offset, dropping at 6dB per 16K bins of offset. The peak spur amplitude depends on ROM width truncation, dropping 6dB per additional address bit used (e.g., a 512x256 addressed ROM -- 512K complex coefficients -- has a peak

spur amplitude of -60dBc, again at ±16K bins offset from the carrier). These results are independent of the FFT computation word size, i.e., identical for 16-bit and 20-bit integer FFTs.

There are several possibilities for calculating the precise coefficients in a truncated ROM; for example, should each entry be the *average* of the multiple "true" coefficients for which that entry substitutes? or perhaps it should be simply the exact coefficient corresponding to the (smaller) FFT for which the ROM is full-sized. We explored this question, trying what we called a "mean ROM" and an "expanded ROM," respectively; we also considered a "median ROM" and a "topographic center ROM". The result of simulation showed that it hardly matters, but where there is a difference the expanded ROM is better. For example, "spurs" of a pure dc input signal are identically zero for the truncated ROM constructed as an expanded ROM, whereas for a mean ROM they are the same size as the carrier spurs (which we might call "ac spurs") described in the previous paragraph, e.g., -54dBdc for a 4M-point FFT using a ROM containing 64K complex coefficients. (These spur amplitudes are for ROMs of 16-bit precision, by the way, whether doing a 16-bit or 20-bit FFT computation. Using instead a truncated ROM of perfect precision has no effect on spur level, which is caused entirely by the ROM's truncated "width.") This perfect suppression of "dc spurs" when using an expanded ROM is less than meets the eye, by the way: when the input data is multiplied by a window function (to reduce spectral "leakage"), as must be done in the real system (see below), the dc spurs reappear, at the canonical level specified in the previous paragraph.

Our next set of simulations involved the addition of uncorrelated Gaussian noise to a discrete array of pure sinusoids, in order to determine how much precision (word size) is needed to ensure that the spectrum of input noise dominates over "numeric noise." This is clearly word-size dependent, since the amplitude of numeric noise in an all-scale FFT equals the LSB (independent of word size), whereas the amplitude of the spectrum of random noise approximately equals the input amplitude reduced by a half bit per butterfly. This estimate suggests that 16 bits is marginal in an all-scale FFT, because even if the input noise level is set to the full-scale amplitude of 2^{15} (a radical approach, allowing no signal headroom), it will emerge in the spectrum at an amplitude of 2^4 after the 22 butterflies of a 4 megapoint FFT; that is +24dB relative to roundoff ("dBr"). The corresponding figure for a 20-bit word size is an output noise amplitude of 2^8 (+48dBr). Of course, one cannot set the input noise amplitude to full scale without severe clipping, owing to the high crest factor of white noise; thus these figures should be reduced by a factor of at least ≈10dB.

The purpose of the noise simulations was to quantify these estimates of the dynamic range, in the output spectrum, of input noise (call it "antenna noise") over roundoff noise ("numeric noise"). We carried out many simulations, with the following result: If the input noise amplitude is set so that approximately 1 sample in 4 million saturates at full scale (call this "full-scale noise"), then the rms noise amplitude in the spectrum that results is approximately 65 (for 20-bit integer arithmetic, with unwindowed input), or 4 (for 16-bit arithmetic). Thus for full-scale noise, the antenna noise in the spectrum has an amplitude +36dBr for a 20-bit computation, +12dBr for a 16-bit computation. In practice one would probably set the input amplitude some 6dB or so below full-scale noise, implying that a 4M-point all-scale FFT must be done at 20-bit precision (or better). The only way to survive with a 16-bit transform is to omit some intermediate scales, a perfectly reasonable (though less conservative) approach.

At this point we decided to use a 20-bit word size, and performed all further simulations at that precision. We next experimented with the "dithering" effects of combining wideband noise with input data, the sum being quantized to 4 or 8 bits. For 8-bit data quantization, dithering increases the dynamic range to ≈90dB (from the 48dB of an undithered 8-bit quantization); however the quantization is a nonlinearity that produces harmonics at the ≈-48dBc level (in addition to the -60dBc spurs caused by a 128K truncated ROM). For 4-bit input quantization the harmonic spurs are far worse, approximately -20dBc, even though dithering continues to provide a wide dynamic range; 4-bit quantization thus appears an unwise choice for SETI.

Finally, we experimented with various window functions. A "window" is jargon for a multiplicative function applied to the input time series for the purpose of reducing sidelobes and leakage: If an FFT is applied to an unwindowed input time series, the finite data length corresponds to multiplication of a continuing time series by a rectangular function (of length equal to the transformed data frame), thus producing in the frequency domain (by the convolution theorem) the convolution of the proper sampled spectrum with a sinc (that's shorthand for sin(x)/x) function, the transform of a rectangle. For an "on-bin" signal (i.e., a sinusoid whose period is an integral submultiple of the transformed time series length) all off-signal bins lie at zeros of the sinc function, producing an accurate spectrum with no sidelobes or leakage; but that is a rare case, and in general one sees sidelobes and signal leakage corrupting the spectrum. The usual cure is to use a multiplicative window function, of unit amplitude at the center of the time series and generally tapering to zero at the ends of the time series (in optics the 2-dimensional analog is known as "apodizing"). The simplest example is the triangle (also called "Bartlett"), but there are literally dozens of contenders for "best window function," named after the famous (and not-so-famous), such as Hanning, Blackman, Dolph-Chebyshev, etc.; for an excellent review see the article by Harris (Proc. IEEE, **66**, 1, 51 [1978]). In general, one trades off improved sidelobe rejection for a broader central response in the frequency domain.

The rectangular window (i.e., no window at all) is a disaster, with peak sidelobe of -13dBc, and slow falloff of sidelobe with offset from the spectral peak (-6dB/octave). At the other extreme, the Blackman-Harris "minimum 4-sample" window has peak sidelobe level of -92dBc, bought at the expense of a factor of ≈2 decrease in spectral resolution (i.e., the response to a pure sinusoid is a peak that spans perhaps 4 or 5 frequency channels before it has fallen off by 30dB). We wished to look at windows because *ii)* we need one, and want to choose rationally, and *ii)* we wanted to see if windowing had side effects on the parameters already simulated (e.g., peak spur level, average noise level, headroom, etc.).

The results are approximately as expected: Windows have negligible effect on spur levels relative to signal amplitudes (because both are similarly affected by the window), etc., and they have the predicted effect on resolution. The average noise level, for "full-scale" input noise, is reduced by about 3dB or so, owing to the reduction of average signal level by the window; this makes the choice of 20-bit arithmetic mandatory, if an all-scale transform is used. Finally, signal amplitudes are reduced by a few dB, relative to numeric noise; this is unimportant, because the system is designed so that antenna noise dominates numeric noise (by some 20dB or more).

The major effect of windowing is to reduce leakage and sidelobes. We tested three windows, namely Hanning (von Hann: a cosine-squared), Blackman-Harris, and triangular (Bartlett), in

comparison with a uniform (no-window) window. Of these, the Blackman-Harris has the lowest peak sidelobe level (-92dBc, falling 6dB per octave offset from the carrier), while the Hanning has only a modest peak sidelobe level (-32dBc) combined however with very rapid falloff away from the carrier (-18dB/octave); the advantage of the Hanning, of course, is a narrower main lobe (1.6 times the width of the uniform window, versus 2.1 for the Blackman-Harris). The choice is not absolute, but depends very much on the nature of the signals and interference expected. For example, if interference is often modulated with audio bandwidths (a few kilohertz), the Blackman-Harris's precipitous drop to -70dBc is of no benefit, and its broader central lobe thus makes it a poorer window. What is needed in this case is a window that confines spectral "splatter" to a handful of contiguous channels, which the Hanning's rapid falloff adequately achieves; thus for this application it is a superior window to the Blackman-Harris because of its superior resolution and sensitivity.

On the other hand, if one is dealing often with interfering *carriers*, the Blackman-Harris is the better window, since it keeps the signal within just a few channels before it falls below the antenna noise continuum. Although the choice is not critical, we believe that experience with the system will dictate which window is better. Thus we are designing the hardware to permit run-time selection -- we will load a suite of windows into the (small) window ROM, selected via a downloaded segment address. Our simulations of windows showed, incidentally, that the window ROM can be truncated enormously with no observable effect: the "full" 4 million coefficients can be replaced by an 8Kx16 expanded ROM (512 times smaller). Thus a single 27C1024 (64Kx16), costing less than ten dollars, can hold 8 window functions.

Based on the simulations just described, we have chosen the parameters for the 4M-point FFT. In particular, we will use 8-bit data quantization (in the mixer-filter-digitizer module), an 8Kx16 expanded window ROM, 20-bit integer arithmetic with all scalers (or all-1) enabled, a full (16Kx16) small twiddle ROM, and a 512x256x16 expanded large twiddle ROM. Figure 13 summarizes the behavior of the FFT with regard to signals, noise, roundoff, and spurs, and Figures 14 and 15 demonstrate the output data from a pair of simulation runs: The "signal" consists of wideband antenna noise to which has been added three large sine waves (amplitude 0.1, at channels 1M, 2M+0.25, and 3M+0.5) and, nearby, three weak sine waves (amplitude 0.001, at channels 1M-8, 2M+20.25, 3M+20.5). In both cases the data has been quantized to 8 bits, and transformed with 20-bit arithmetic in an all-scale FFT using a 512x256x16 expanded large twiddle ROM. The output table prints complex pairs, 4 to a line, beginning at the labelled channel number. In Figure 14 we have used a uniform window: the on-bin signal at 1M is cleanly resolved (one bin, all real), allowing clear detection of the nearby signal (at 1M-8); but the off-bin signals (at 2M+0.25 and 3M+0.5) are broadened by "spectral leakage" to more than the 68 contiguous channels shown, burying the nearby weak signals. In Figure 15 we have used the same data and transform parameters, but with a Blackman-Harris window (truncated to 16Kx16). Now the on-bin signal at 1M has been broadened to a half dozen channels, somewhat degrading its detection, but the weak signal at 1M-8 is still cleanly resolved. More importantly, the off-bin signals at 2M+0.25 and 3M+0.5 are of comparable width, dropping below the antenna noise level at ±3 bins -- the nearby weak signals (at 2M+20.25, 3M+20.5) now show clearly! Given that only a small fraction of real-world signals are on-bin, the wisdom of windowing should be apparent.

In these tables, the small peak near 4M is a harmonic spur of the strong signal at 2M, at a level a few times the average antenna noise level. The spectral regions offset by 16K from the large peaks have been listed because the worst ROM-truncation artifact occurs there; even with these relatively strong signals the spur does not rise above antenna noise (it is -60dBc, corresponding to an amplitude of about 15 in Figure 15).

2.4.2 FFT Circuit Implementation

We have delayed laying out a printed circuit board for the FFT processor until we see actual A41102 silicon, preferring to concentrate our efforts on the upstream and downstream modules (antenna, downconverter, hardware backend, etc.). We have, however, designed the processor at the block diagram level. It follows the description above, with "glue" logic implemented in complex PLDs (AMD Mach series). Each 4M-point FFT requires 9 4-megabyte parity DRAM SIMMs, a handful of buffers and latches, the twiddle and window ROMs, and the PLD glue. We have included a diagnostic bus for debugging and maintenance; it exposes the data stream as it propagates through the circuit. We have devised an interesting quasi-logarithmic compression scheme for converting quadrature amplitude pairs into modulus; we have not seen it elsewhere, but believe it is optimum for constant fractional error compression of integers. We will describe it in the next semi-annual progress report.

On the chip supply issue (a continuing saga): We placed an order for 250 of these chips in December, 1991; subsequently Austek's rights to the chip (granted by the governmental CSIRO of Australia, which funded development of the chip) expired, with supply picked up by the large Australian AWA Corporation. Although they projected delivery for October, 1992, we still have not received any parts.

2.5 Backend Array

The FFT array is followed by an array of dedicated digital signal processing circuits, which sift through the spectra looking for *i)* new peaks that exceed background by a significant amount, and *ii)* the progress of previously flagged frequencies of interest; see the original proposal for a detailed description.

We had assumed that this function is best performed by an integrated digital signal processor (DSP). However, after looking at several choices for the DSP array, including the Motorola 56002, 96002, Intel 80960, DEC Alpha, and Star chip, we have decided to use a hybrid approach -- a *hardware* baseline accumulator and thresholding circuit, followed by a general-purpose computer array for decision-making. The hardware approach provides the same performance as a pair of DEC Alpha chips, currently the fastest available general-purpose processor, at a small fraction of the cost (we estimate our module will cost about $350 per board). Using hardware, of course, generally eliminates flexibility; we hope to recover it here by the use of downloaded search parameters (see below).

Figure 16 is a block diagram of the hardware baseline module. It accepts 16-bit spectral moduli from the FFT processor (computed as described in the previous progress report), keeping a running boxcar average of channel power via a SRAM-based circular buffer. It also accepts a

download of a single threshold constant, followed by a table of signal frequencies to watch (a "slot") or ignore (a "notch"), via the FIFO at lower right. The hardware compares the running average with the current datum (via the downloaded threshold), sending back the frequency, modulus, and baseline constant of all peaks that exceed the designated threshold. At the same time it also forwards the modulus information for the frequencies in the FIFO slot table, and ignores those in the notch table. Slots and notches are specified by beginning and end, so they can be arbitrarily wide in frequency (e.g., modulated interference).

We have simulated the hardware module, including the quasi-logarithmic compression algorithm, and it works perfectly. We intend to reduce this module to hardware during the current 6-month grant period.

2.6 Backend and Workstation

The hardware threshold modules communicate with an array of general purpose processors. After looking at various options (an array of VME single-board computers, a single very fast computer, etc.) we have decided to use an array of inexpensive PC-clone motherboards. These are commodity items, costing roughly $500 for a bare motherboard with a 486SX/25 processor. The performance of commodity motherboards (10 to 15 MIPS) is comparable to or better than special purpose VME solutions, at around a third the cost. The use of PC-clone motherboards also allows us to use inexpensive, widely available software tools instead of a high cost, specialized software development platform.

We intend to use twenty motherboards, networked with thin-wire Ethernet, in a diskless configuration. Each will boot and receive startup instructions over the network from the controlling workstation. Because this eliminates the need for local disk drives, hardware reliability is improved and software version control problems are eliminated. All network communication will be done with straight Ethernet protocols. We looked at using higher level, error tolerant methods such as TCP/IP and other Internet protocols. Since our application is relatively straightforward and the network hardware is in a very simple configuration (a single Ethernet cable), we found that these high level protocols gave us very little benefit while requiring a lot of overhead.

Having an Ethernet full of computers running nearly identical code in lockstep might seem to be a recipe for disaster because of the high probability of packet collisions. This problem will be almost eliminated in our application since the workstation will rigidly control network communication. It will normally poll the diskless motherboards and only allow them to transmit onto the Ethernet at specified times. Unsolicited transmitting will only occur at boot up time and under error or other rare event conditions when latencies caused by Ethernet's "random delay/retry" policy can be tolerated.

We are in the process of writing a custom boot ROM to be used in the ethernet cards which will be plugged into the PC motherboards. Commercially available boot ROMs perform a similar function, but only work with proprietary PC networking software. Our ROM's function is simple: it will make a file on the workstation look like a bootable floppy disk for each PC motherboard. The motherboard can then boot MS-DOS and run normal DOS programs. Our

specialized processing software can therefore be developed on a PC using common software tools and techniques.

The software running on the motherboards will accept the baseline and hit data from three of the hardware baseline modules (3 feedhorns, a single 2 MHz sub-band), which will be implemented as ISA plug-in boards. It will then correlate these pieces of data with each other, with past data, and with knowledge about the antenna's beam profile. It will reject hits that are wide enough to be natural or obviously of terrestrial origin. Hits in either of the sky horns that coincide with strong hits in the terrestrial horn will be rejected as obvious terrestrial interference. Hits that traverse the east sky horn and then the west sky horn, with proper amplitudes, will be flagged as possible candidate signals and this information will be forwarded to the workstation for further processing and archival storage.

We are investigating a simple algorithm that analyzes data from the two sky horns to give us a measure of how closely a set of hits conforms to a possible extraterrestrial signal. It takes advantage of the overlapping beams to remove ambiguities caused by fluctuating signal strengths (from interstellar scintillation or polarization shifts). The algorithm takes the difference in signal strengths between the two sky horns and normalizes it by dividing by their sum. When plotted versus time, this normalized difference will follow a specific path, which is a function of the beam shapes and the earth's sidereal rotation. If a set of candidate hits does not follow this path, it can be rejected since the transmitter is not moving sidereally with respect to the earth.

We have made some progress in designing algorithms for the backend workstation. In particular we have written an expandable shell for a graphical user interface, and we are experimenting with algorithms for handling the spectral data. We have a rudimentary simulator that creates the kind of data that will eventually come from the spectrum analyzer, thereby allowing us to test these algorithms and their associated display windows.

3. NEXT STEPS

We have made major progress during this 6-month period, on the antenna, amplifiers, downconverter, digitizer, FFT architecture and simulations, hardware threshold module, and backend. The first three are complete or nearly so, and we expect to complete the digitizer hardware soon. The FFT array and following subsystems constitute a major construction effort, particularly with our expanded goals (240 million channels and three antenna feeds, compared with the 100 million originally proposed), and we expect those tasks to occupy a full year's effort. Overall we are pleased with progress to date.

4. OTHER FUNDING

During the period of this report we have received funding from The Planetary Society and from the Bosack/Kruger Charitable Foundation, in addition to our grant of partial support from NASA; we also received equipment donations from the John Fluke Company and the Hewlett Packard Corporation.

5. PUBLICATIONS AND TALKS

The two papers listed in the previous progress report are in press and awaiting acceptance, respectively.

During the period of this report we have given talks on the project at Boston's Museum of Science, the ARRL Convention at Boxboro, MA, and the ASC Conference at MIT. Our work has been filmed and shown on CNN, local TV channels 5, 7, and 56, British channel 4 (ITV), and NPR radio.

6. ACKNOWLEDGEMENTS

We are grateful to Michael Davis (Arecibo Observatory), Jon Hagen (NAIC), John Kraus (Ohio State University), David Staelin (MIT), and Dan Werthimer (UC Berkeley) for valuable discussions. Martin Gimersky (University of Victoria) provided major assistance with the antenna pattern calculations.

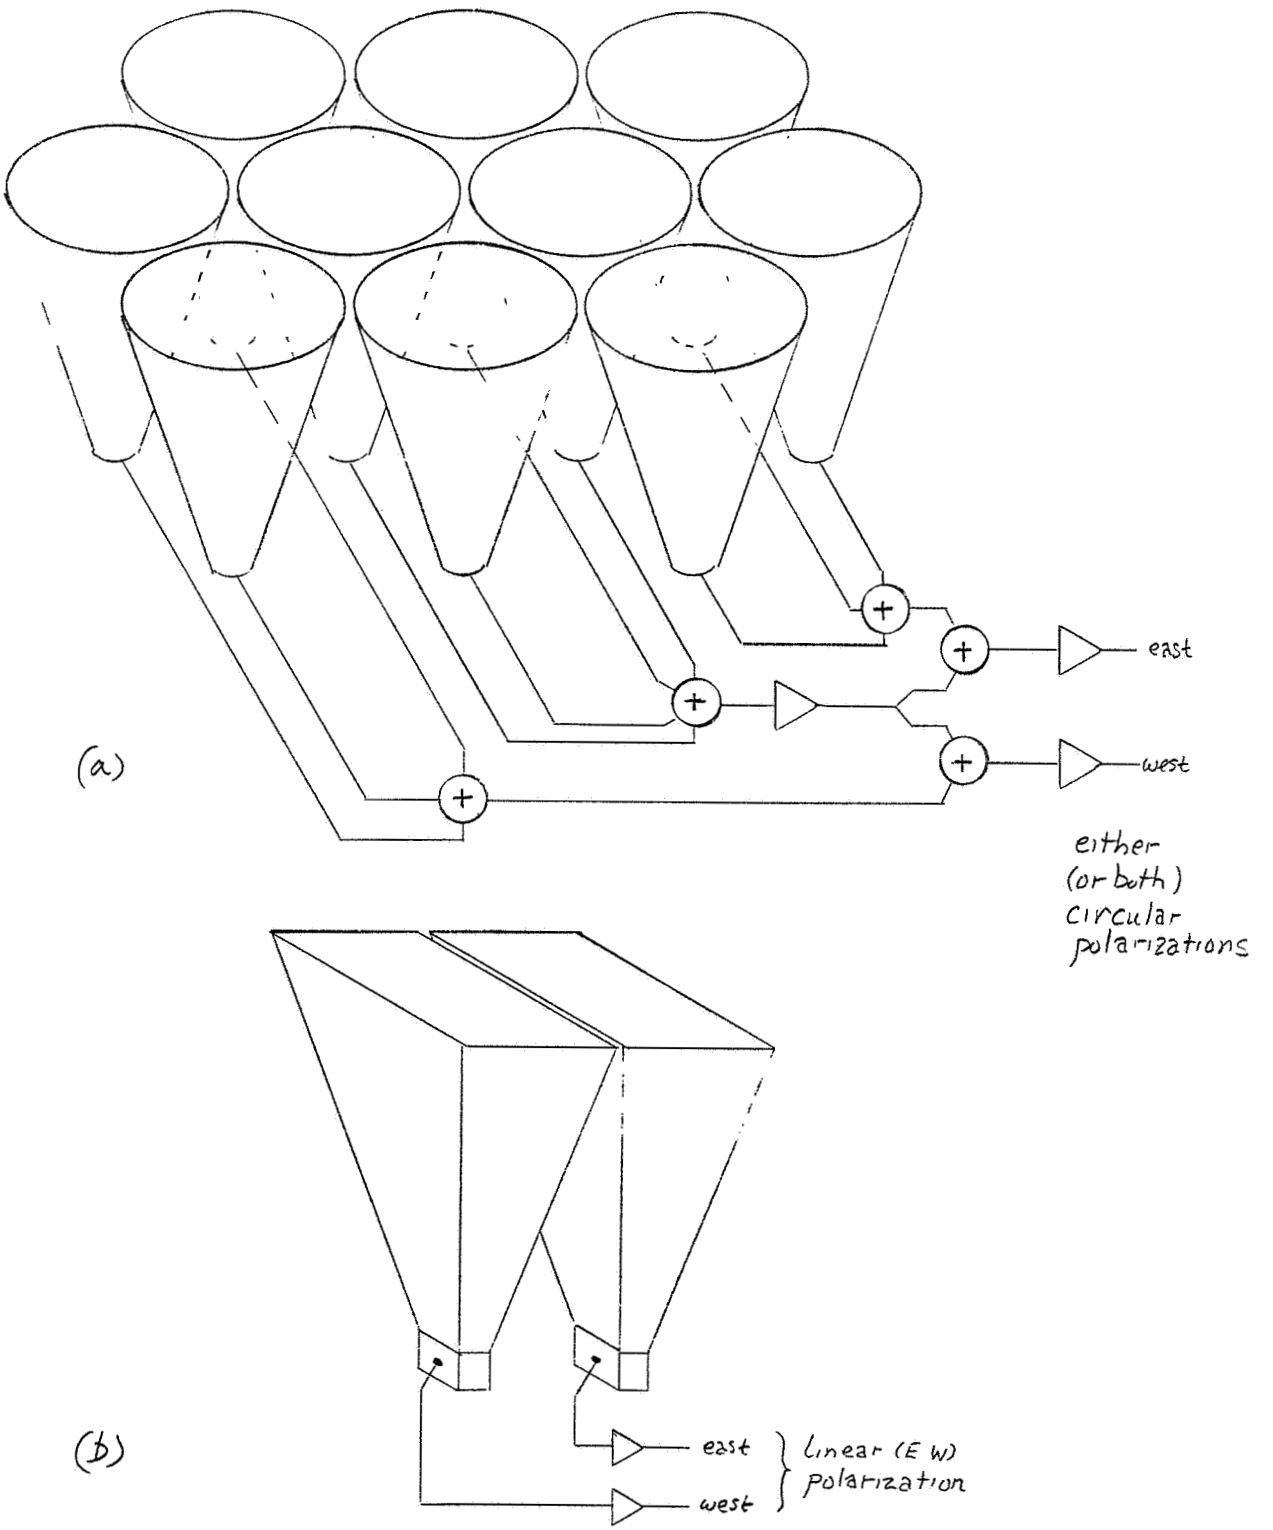

Figure 1. Dual-beam focal plane alternatives: (a) phased array of dual-polarization circular horns, using passive combiners and low-noise amplifiers (only one polarization shown); (b) pyramidal linear-polarization horn pair, stacked along the E-plane.

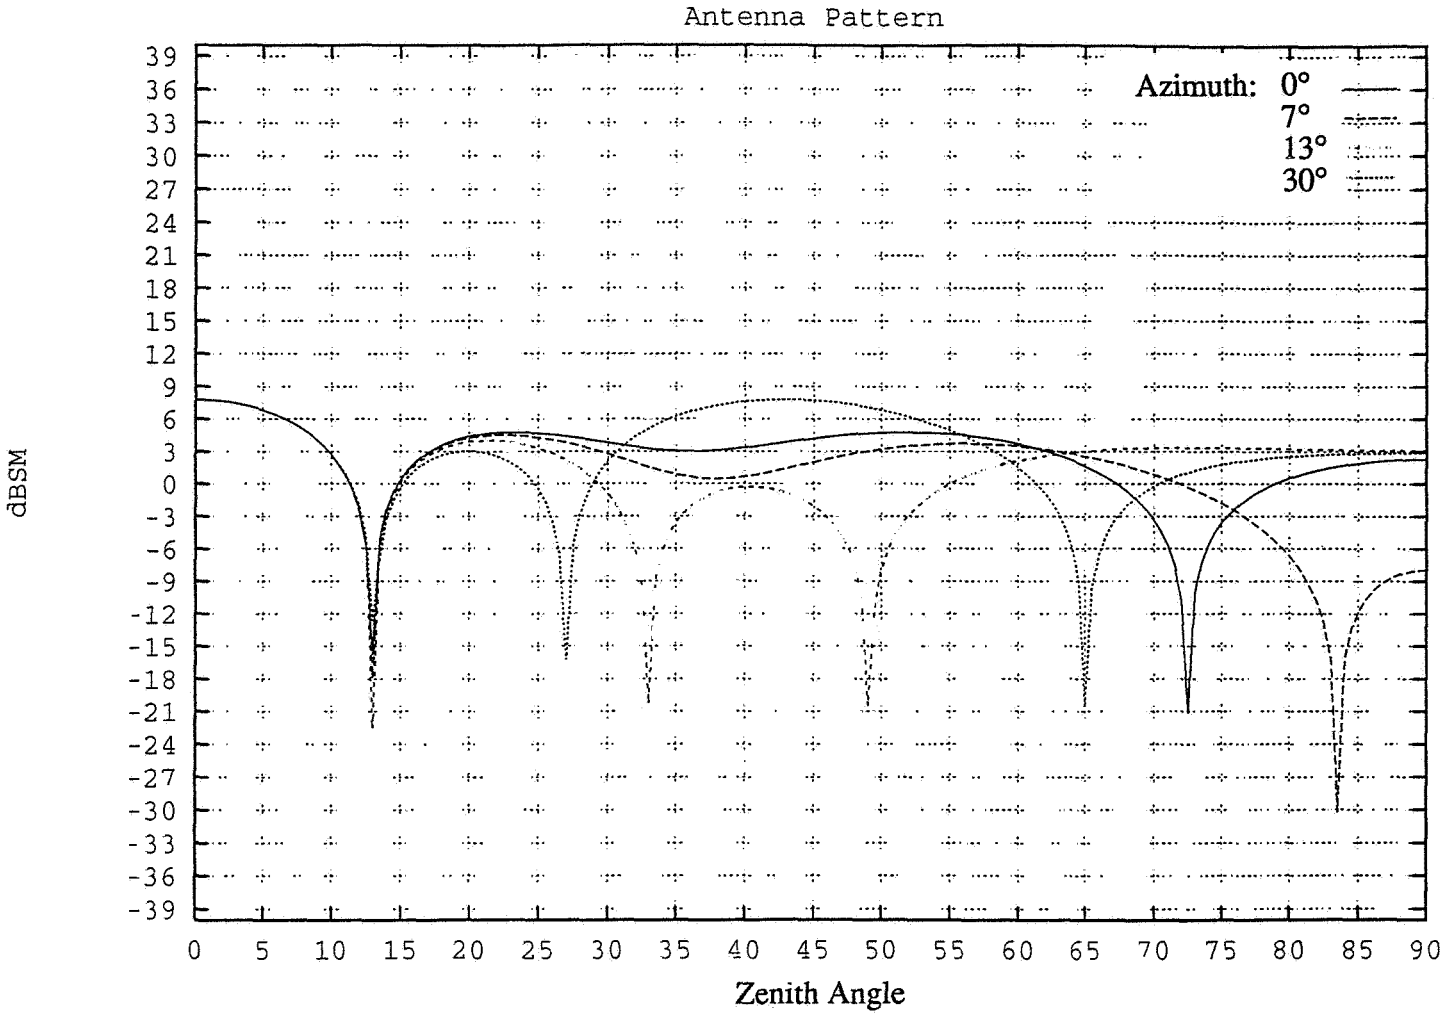

Figure 2. Antenna gain as a function of zenith angle (for four choices of azimuth), for a hexagonal phased subarray of non-interacting point radiators, with lattice-plane spacing of 1.5 wavelengths. For an array of finite sized conical feeds, as in Figure 1a, this plot must be multiplied by the single-horn diffraction pattern, which largely suppresses the off-axis grating lobes. Note that the use of dual circular polarization feeds, with separate combining networks, allows one to construct a dual-beam, dual-polarization focal plane array. Though the pattern shown is for a hexagonal array, the result for a heptagonal array is nearly identical. The close matching of main-lobe patterns is maintained for all azimuth angles; the particular choice here was meant to be "random."

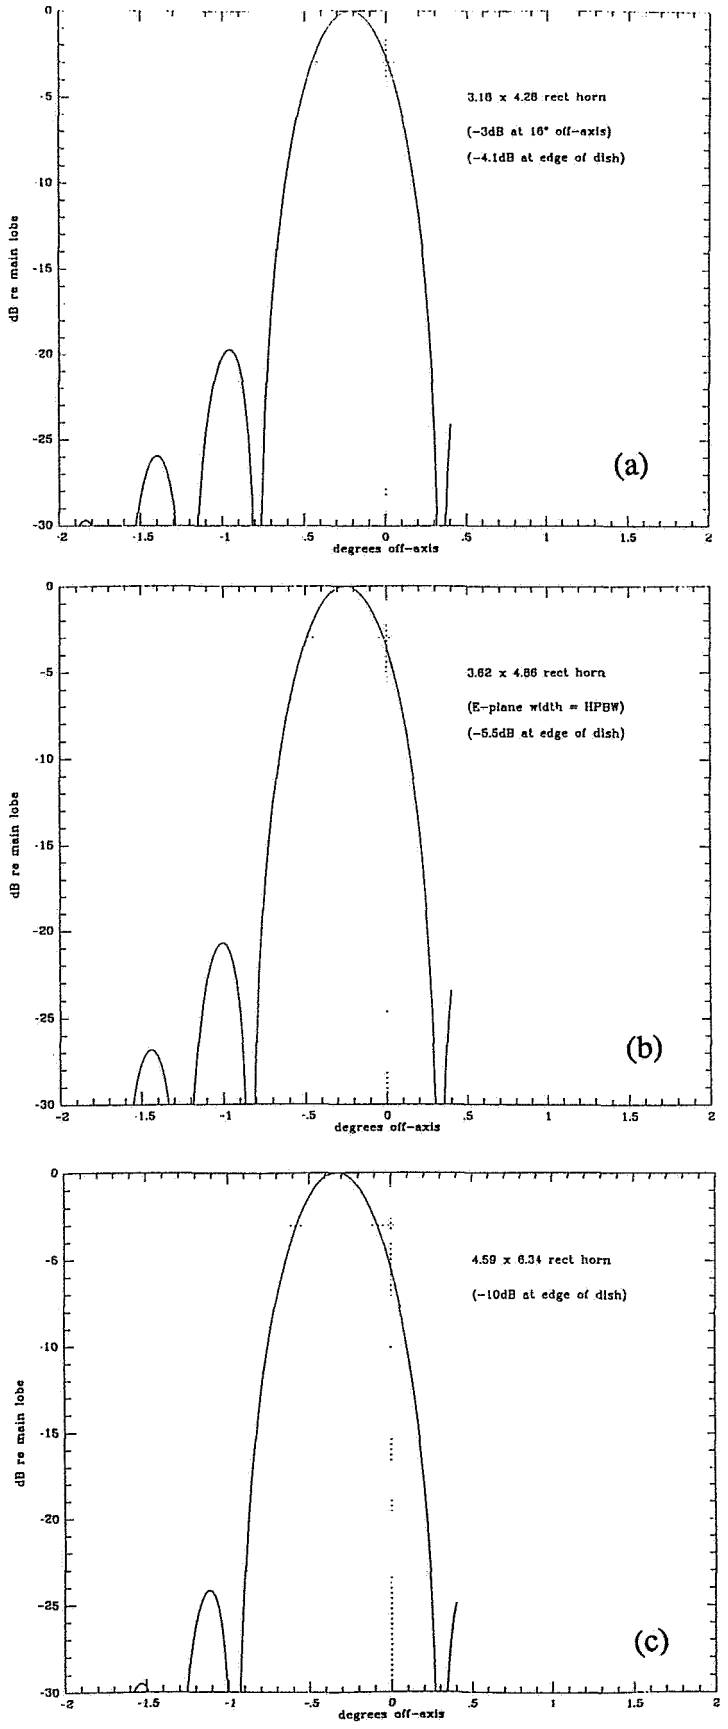

Figure 3. Far-field antenna pattern for a 26-meter Cassegrain illuminated by a single pyramidal horn that is displaced along its E-plane by half its aperture. The three cases plotted progress to larger apertures, specified in wavelengths, with edge tapers of -4.1dB, -5.5dB, and -10dB, respectively. Each graph is centered on the H-plane, and plotted versus angle in the E-plane, with the vertical dashed line indicating the antenna axis, and the horizontal dashed line indicating -3dB relative to maximum gain. The horns are assumed non-interacting; the plots were terminated at +0.4° because of memory limitations.

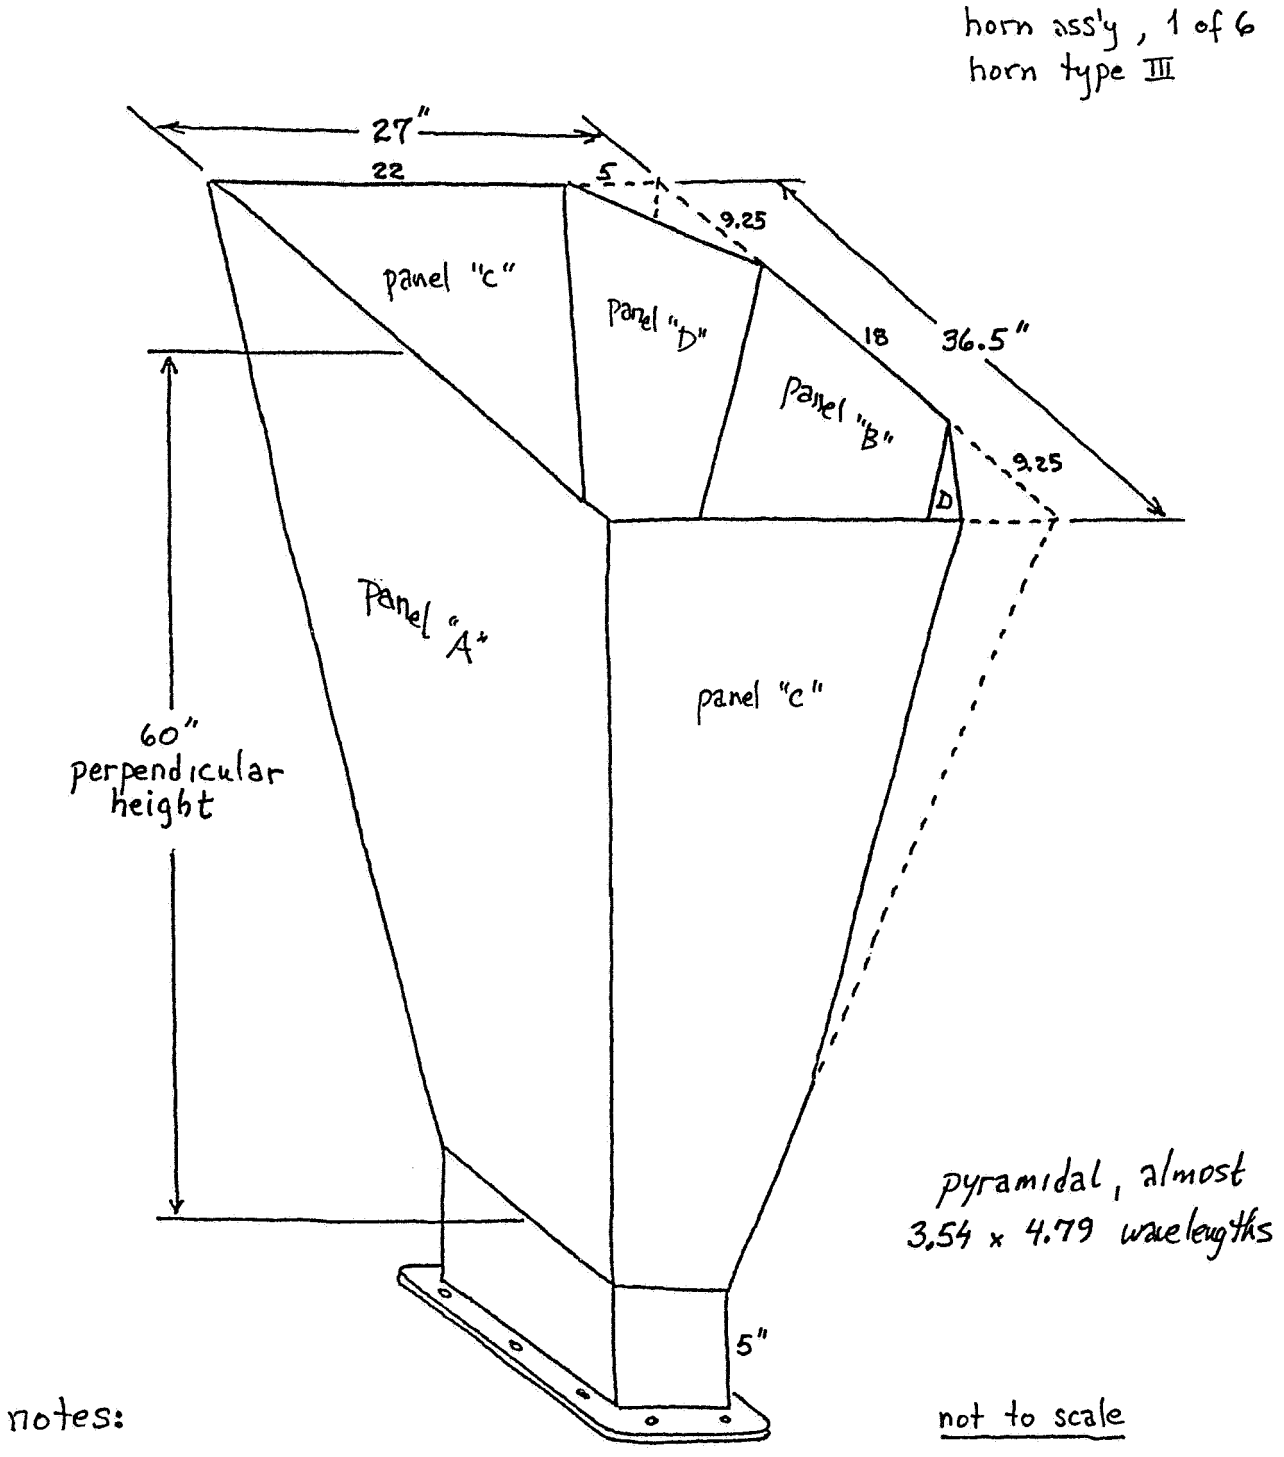

Figure 4. Pyramidal horn design, with truncated corners to allow placement in our radome. This horn is closest to Figure 3b.

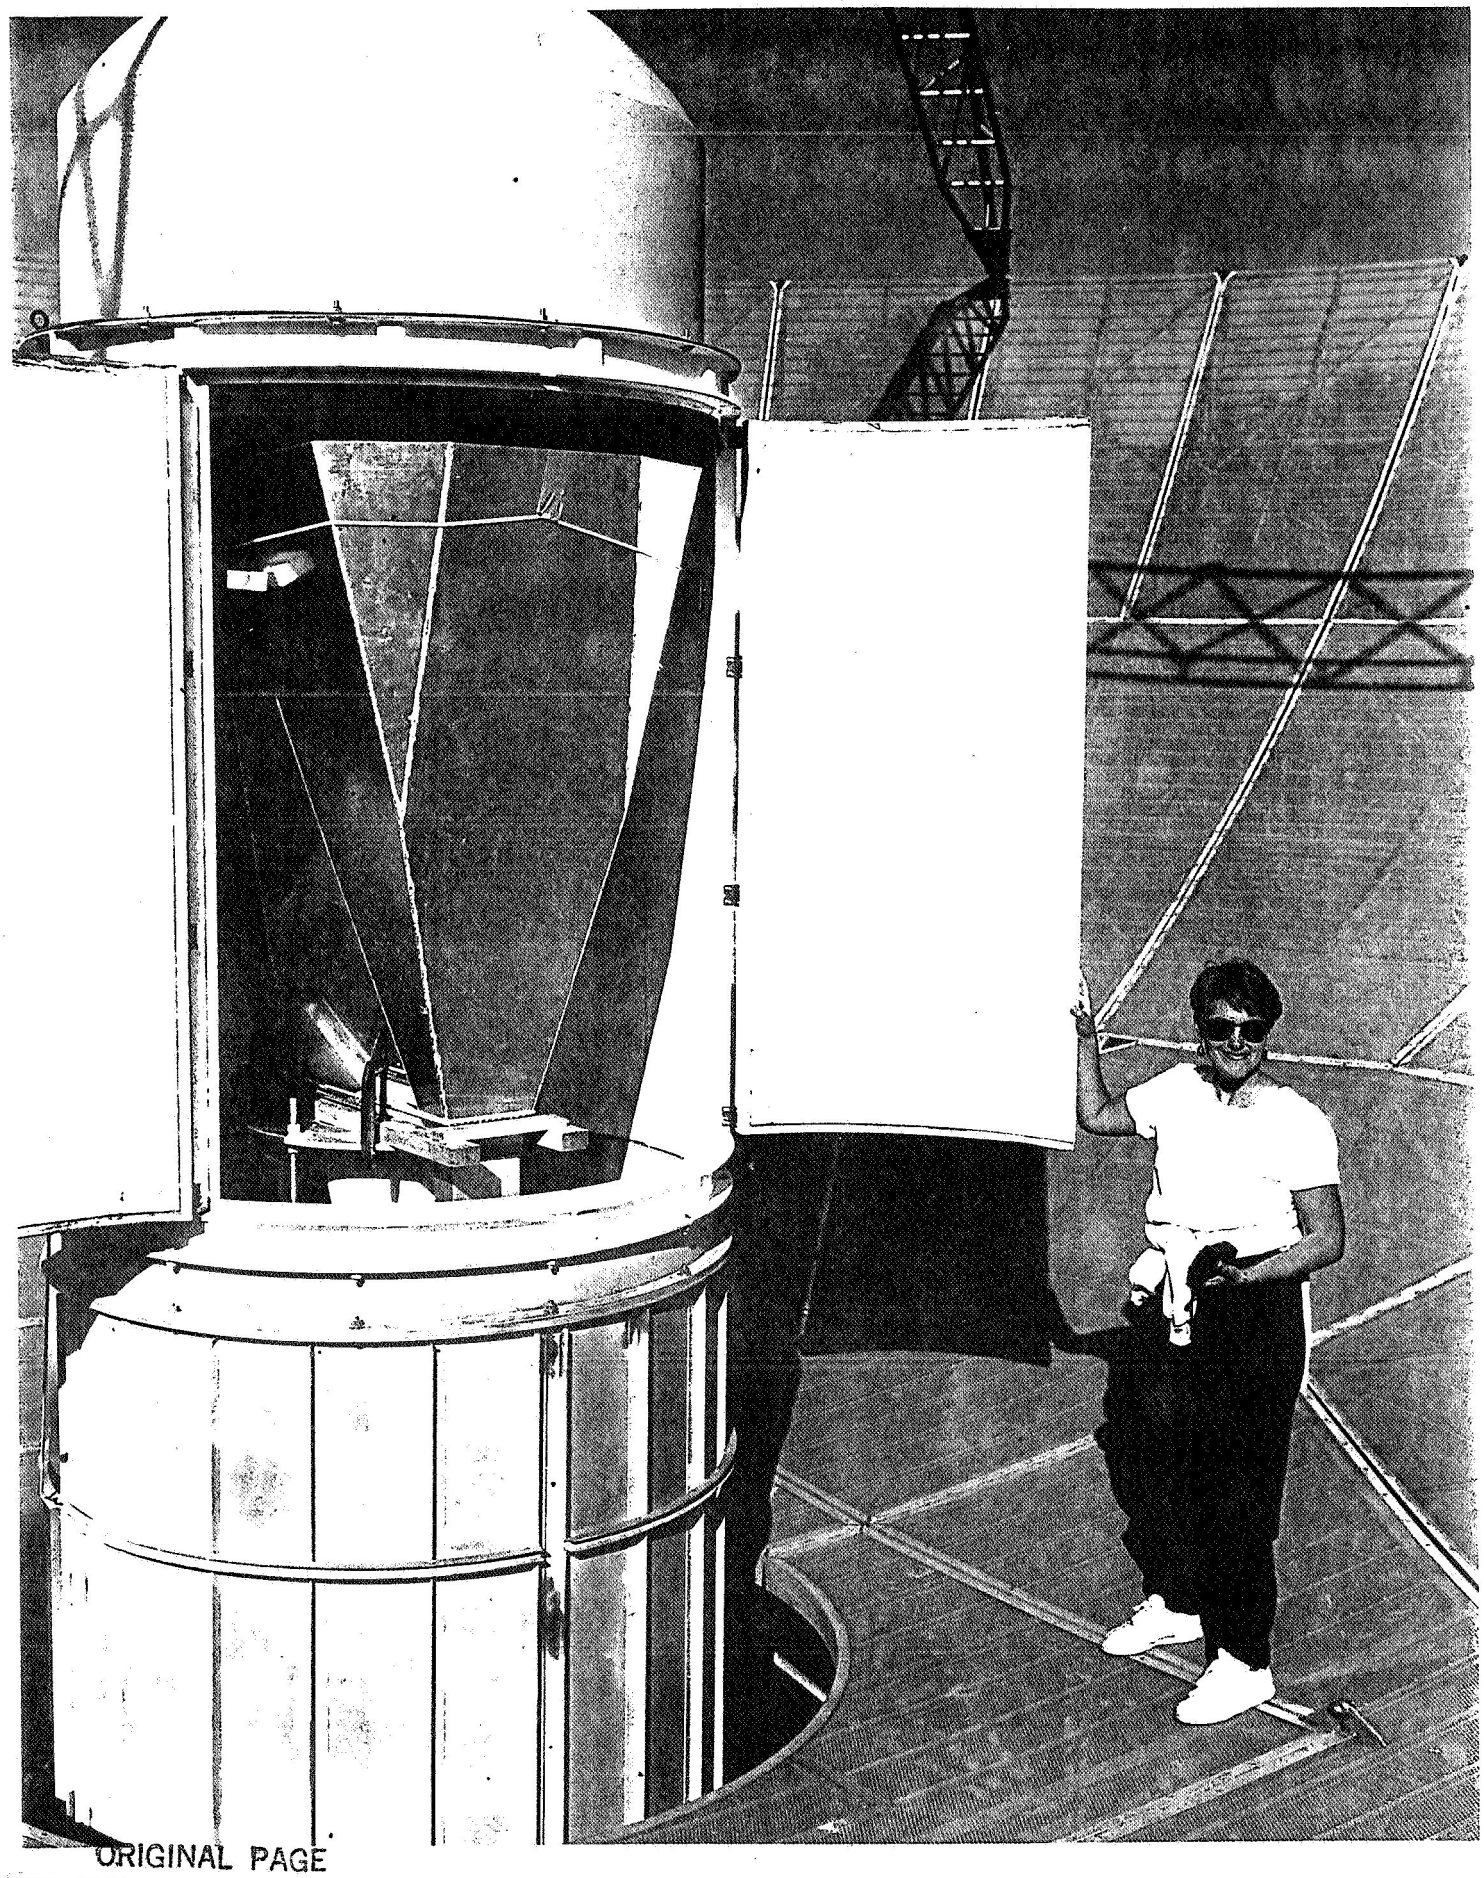

ORIGINAL PAGE
BLACK AND WHITE PHOTOGRAPH

Figure 5. Dual pyramidal horns, fabricated as in Figure 4, mounted in the 26-meter Cassegrain radome. Because of mechanical constraints, the phase centers of the horns are approximately 60 cm too close to the subreflector; this can be corrected when warm weather returns. A portion of the subreflector support truss is visible at top.

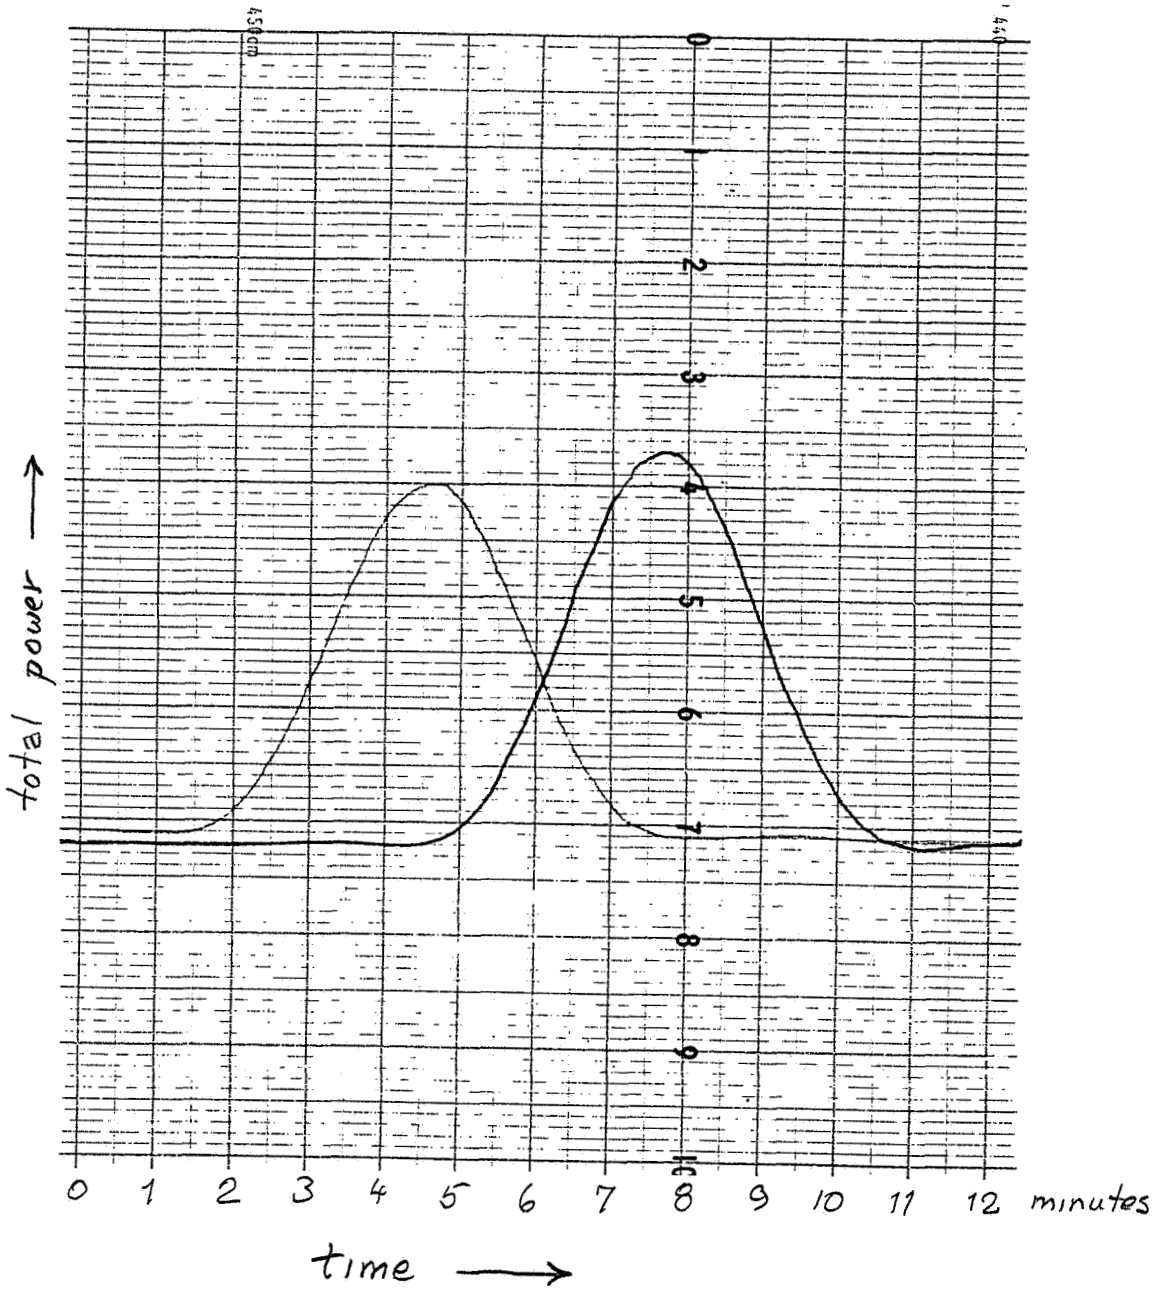

Figure 6. Sidereal drift scan of Cyg A (declination≈40.6°) with dual pyramidal horns, showing good lobe pattern and handoff, but lower than expected signal strength. The true beam separation is 0.5 major divisions larger than shown, due to pen offset in the strip chart recorder.

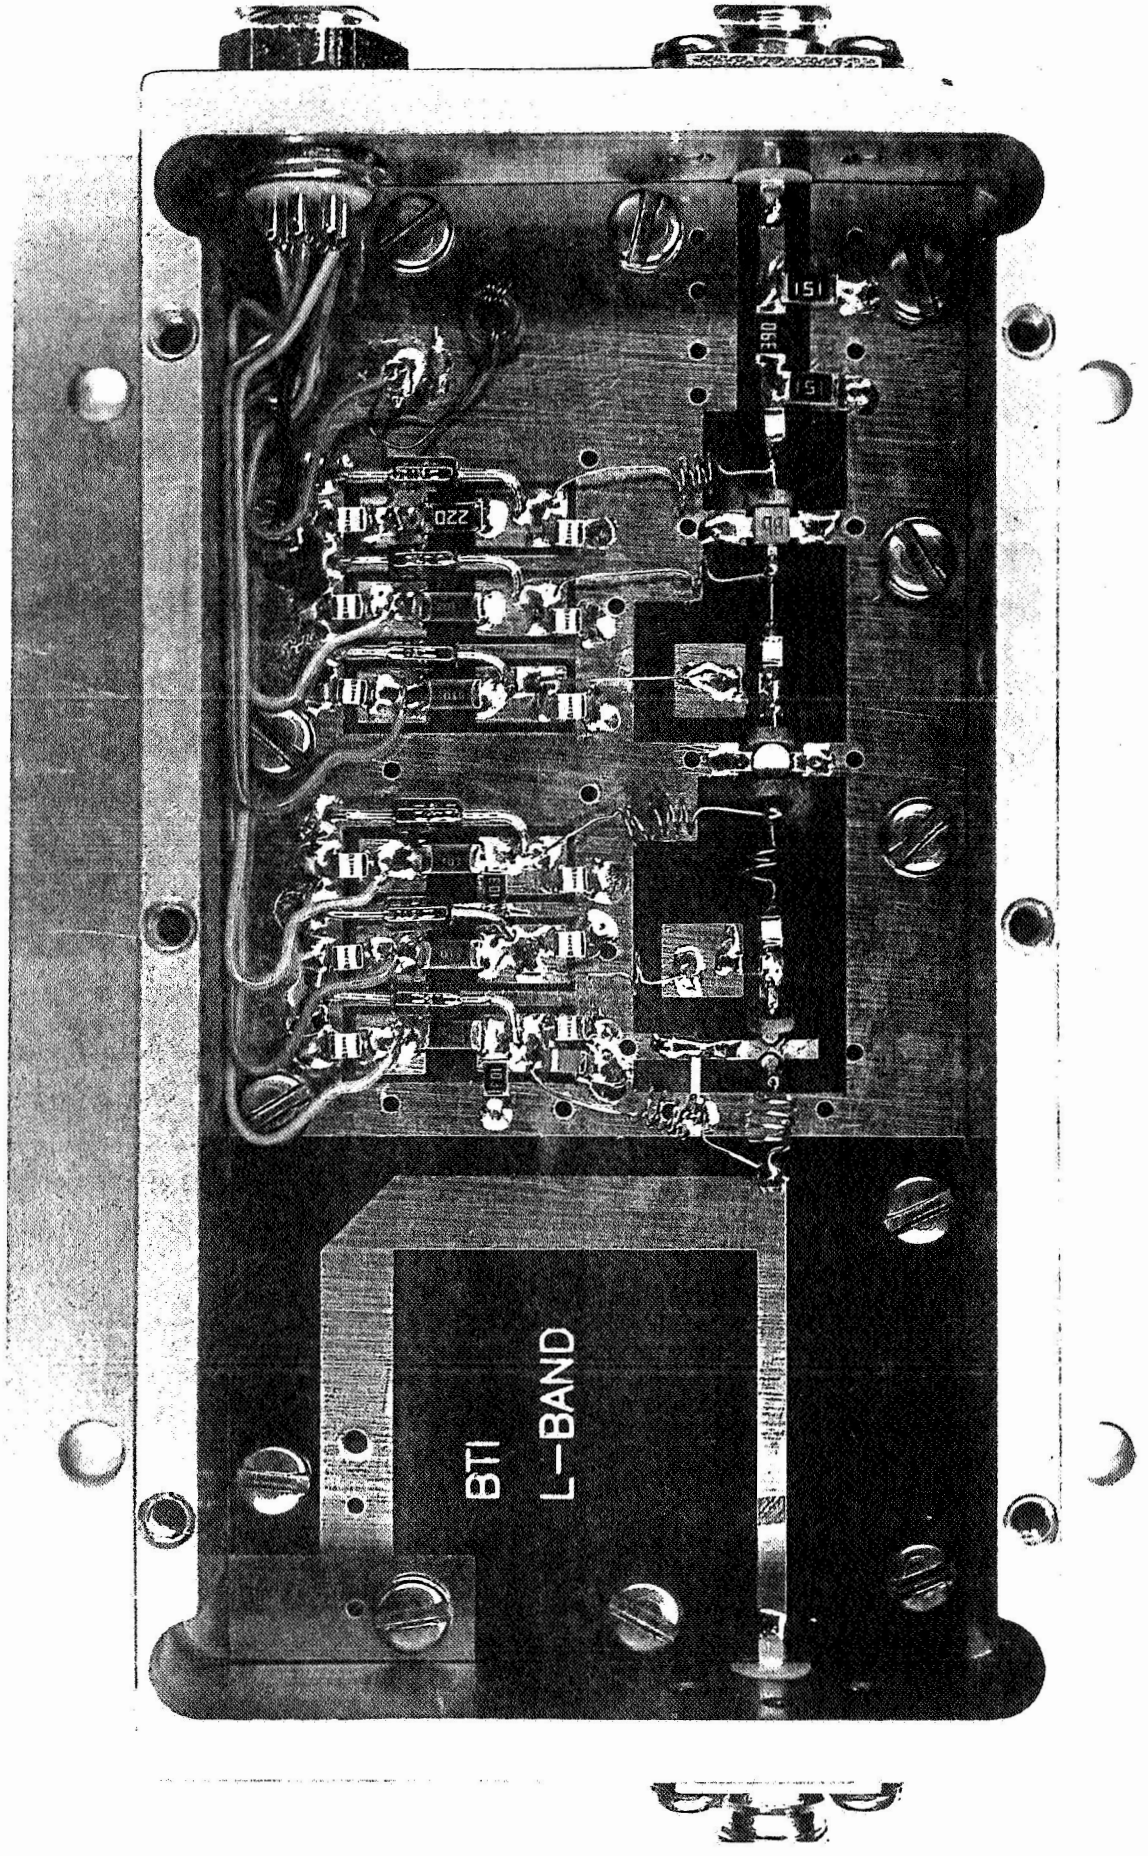

Figure 7. 3-stage HEMT low-noise broadband amplifier, constructed from a kit of parts supplied by Berkshire Technologies. The photograph is considerably enlarged; actual amplifier dimensions are 1.75"x3.25".

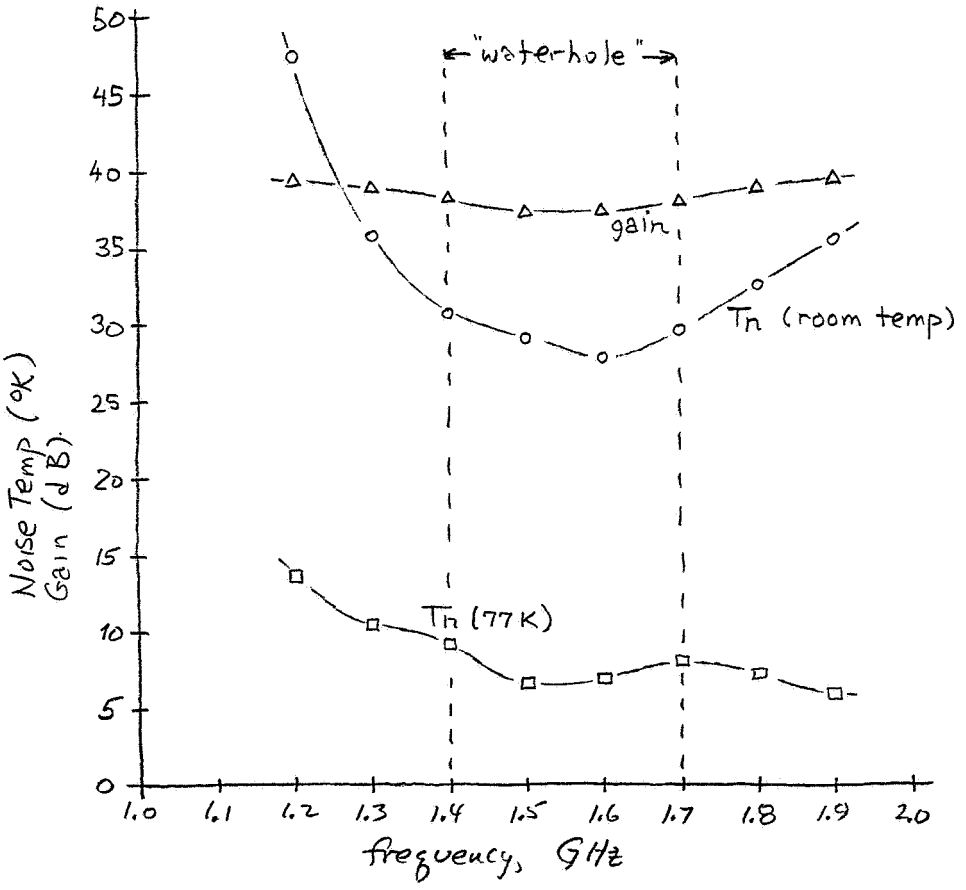

Figure 8. Noise and gain of the HEMT amplifier shown in Figure 7.

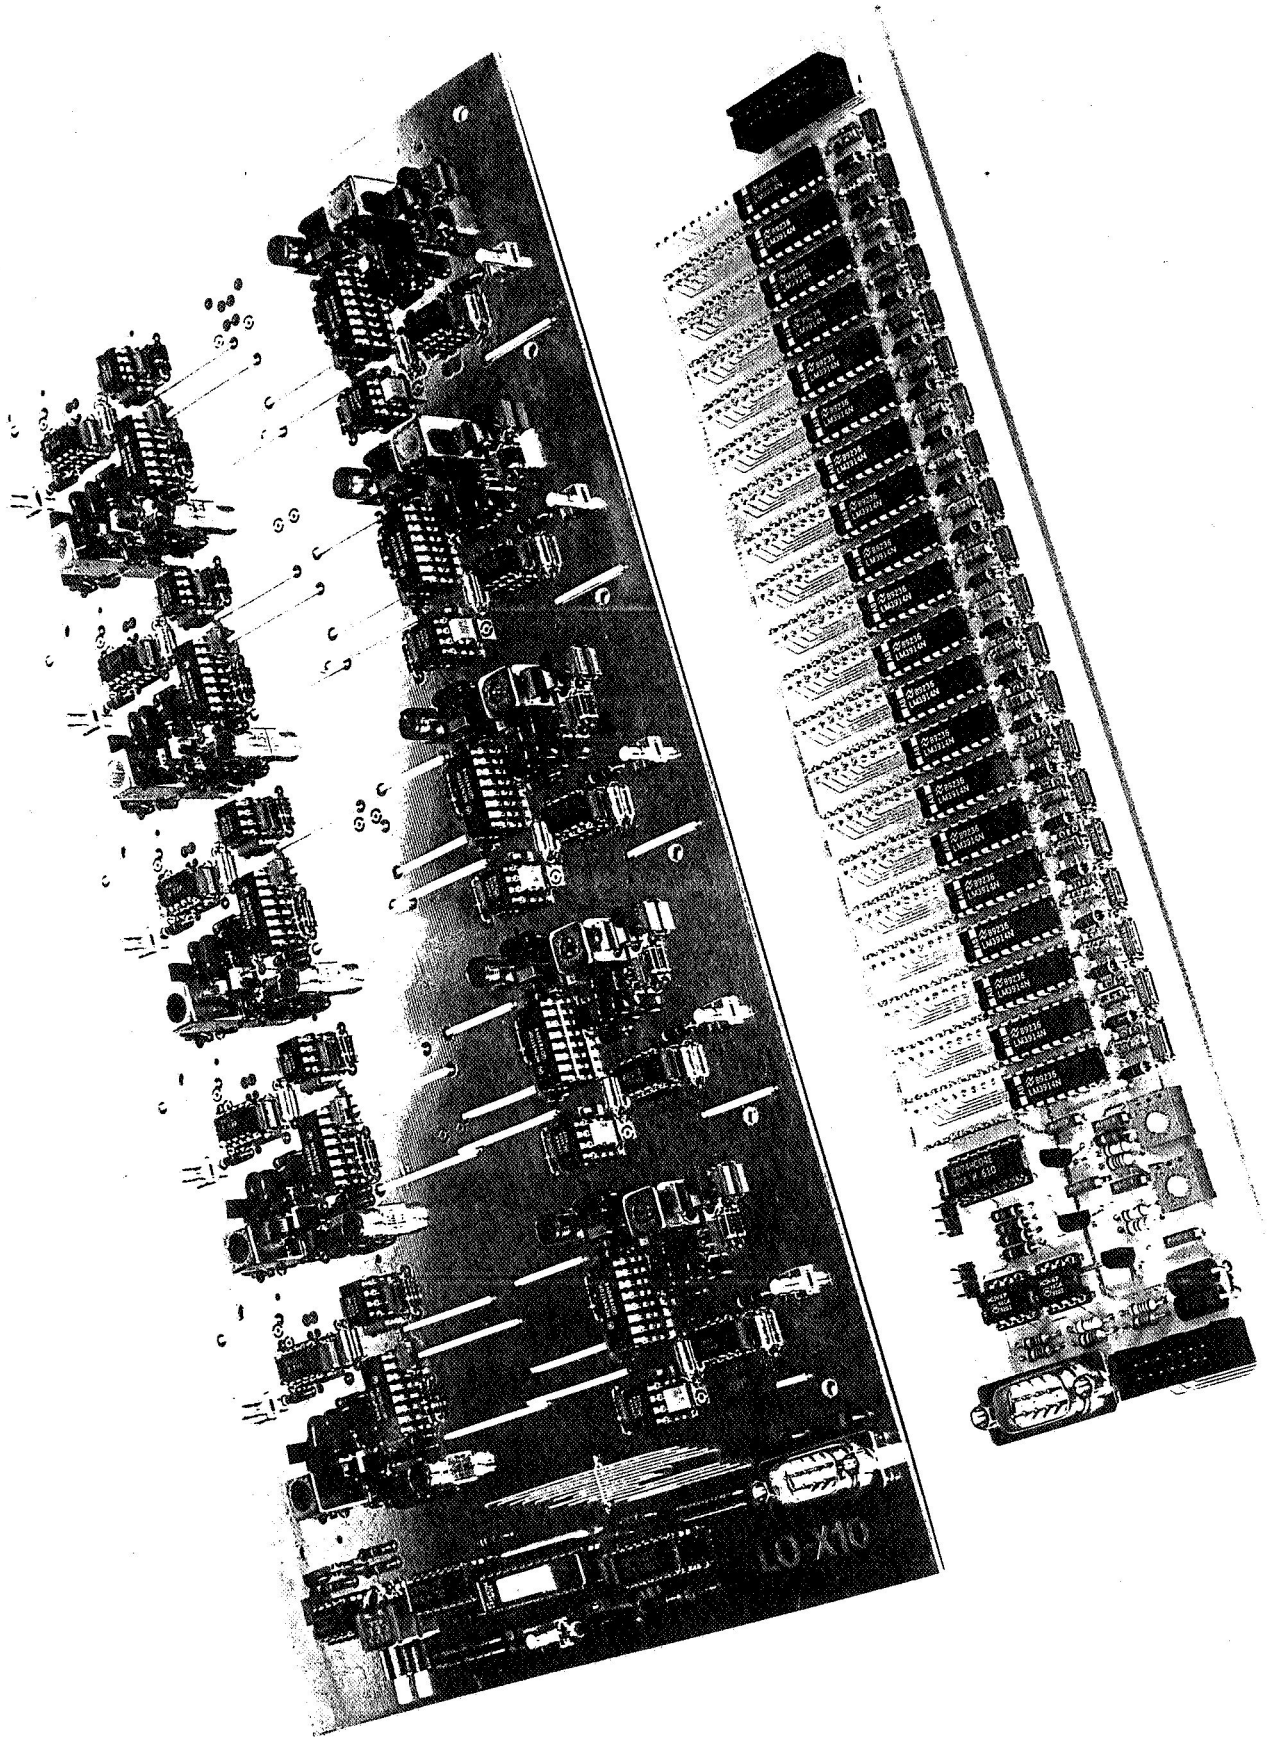

Figure 9. Completed VHF local-oscillator array and bargraph display boards. The LO synthesizers receive their serial register load from an 87C751 microcontroller, which then enters the irreversible "power-down" mode in order to prevent digital noise. Two synthesizer boards provide the full 20 oscillator array for the 240-megachannel spectrometer.

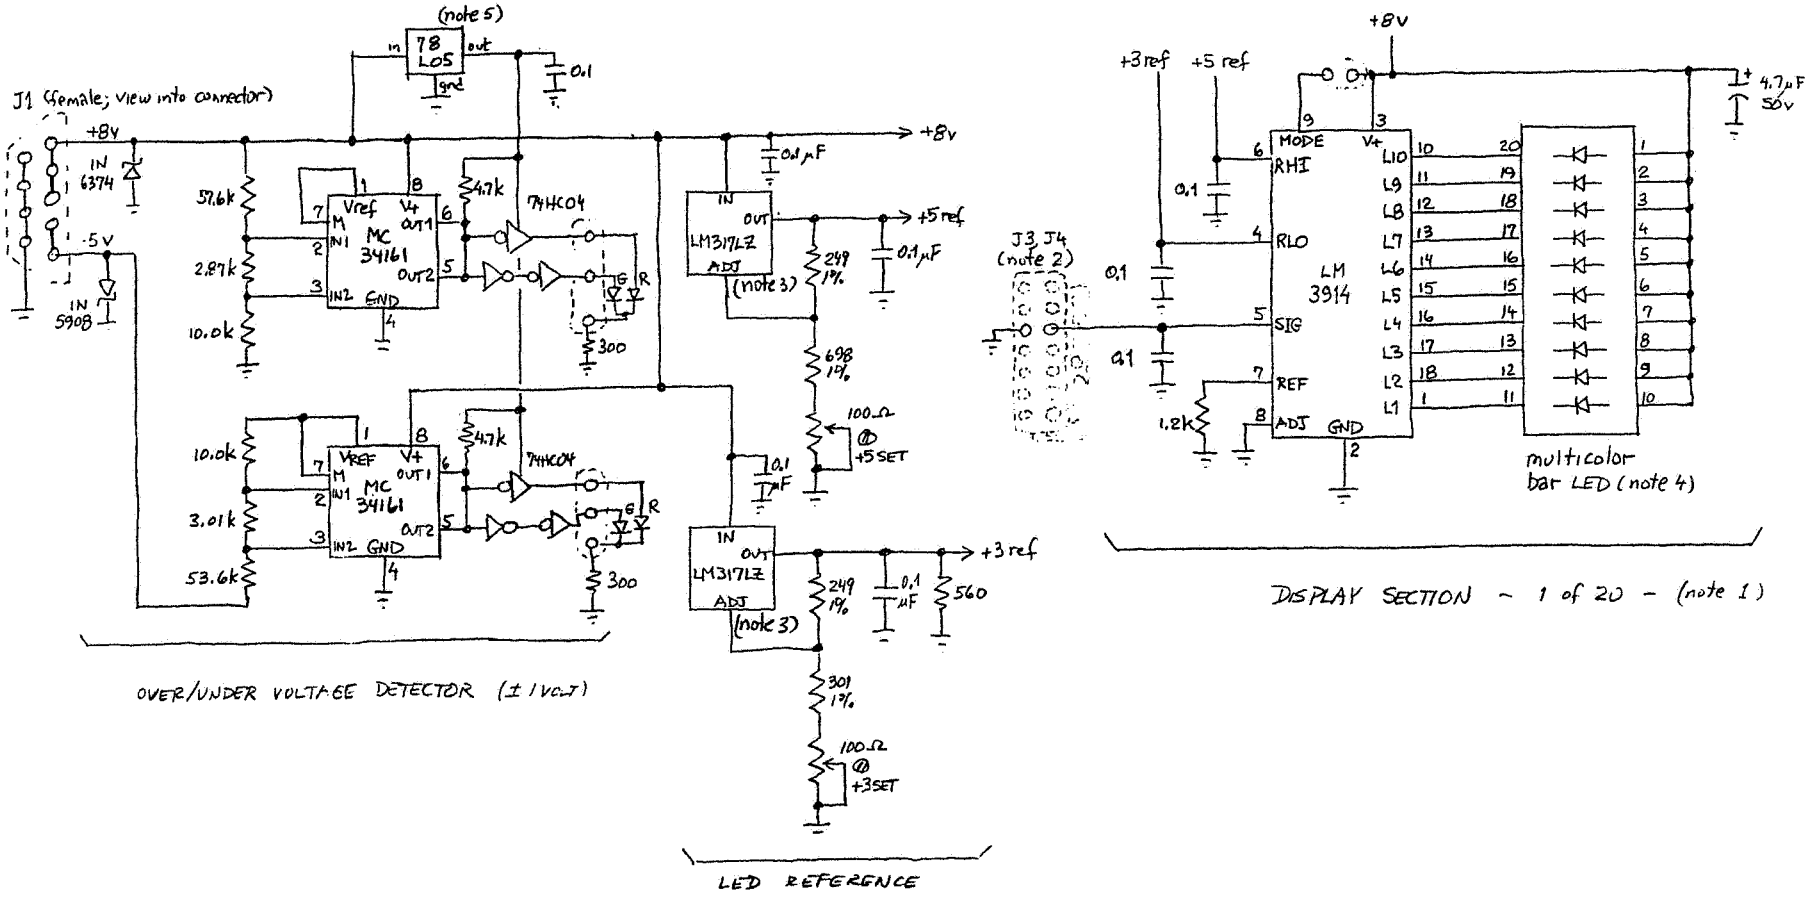

Notes:

(1) there are 20 LED bars, one for each PLL. each must have the bypass caps shown

(2) PLL levels arrive on 14-pin, 2-row headers, each header carries 10 of the 20 PLL levels pinout is as shown:

(3) TO-92 package:

(4) LM3914 must be placed vertically, pin 1 down, with LED bar on high-numbered side, that is because L1 is lowest level, L10 is highest

(5) TO-92 package:

PLL DISPLAY & VOLTAGE MON.		
rev	date	comments
0	11/15/92	orig drwg

Figure 10. Schematic of bargraph display board, which also includes a power-supply over/under voltage indication.

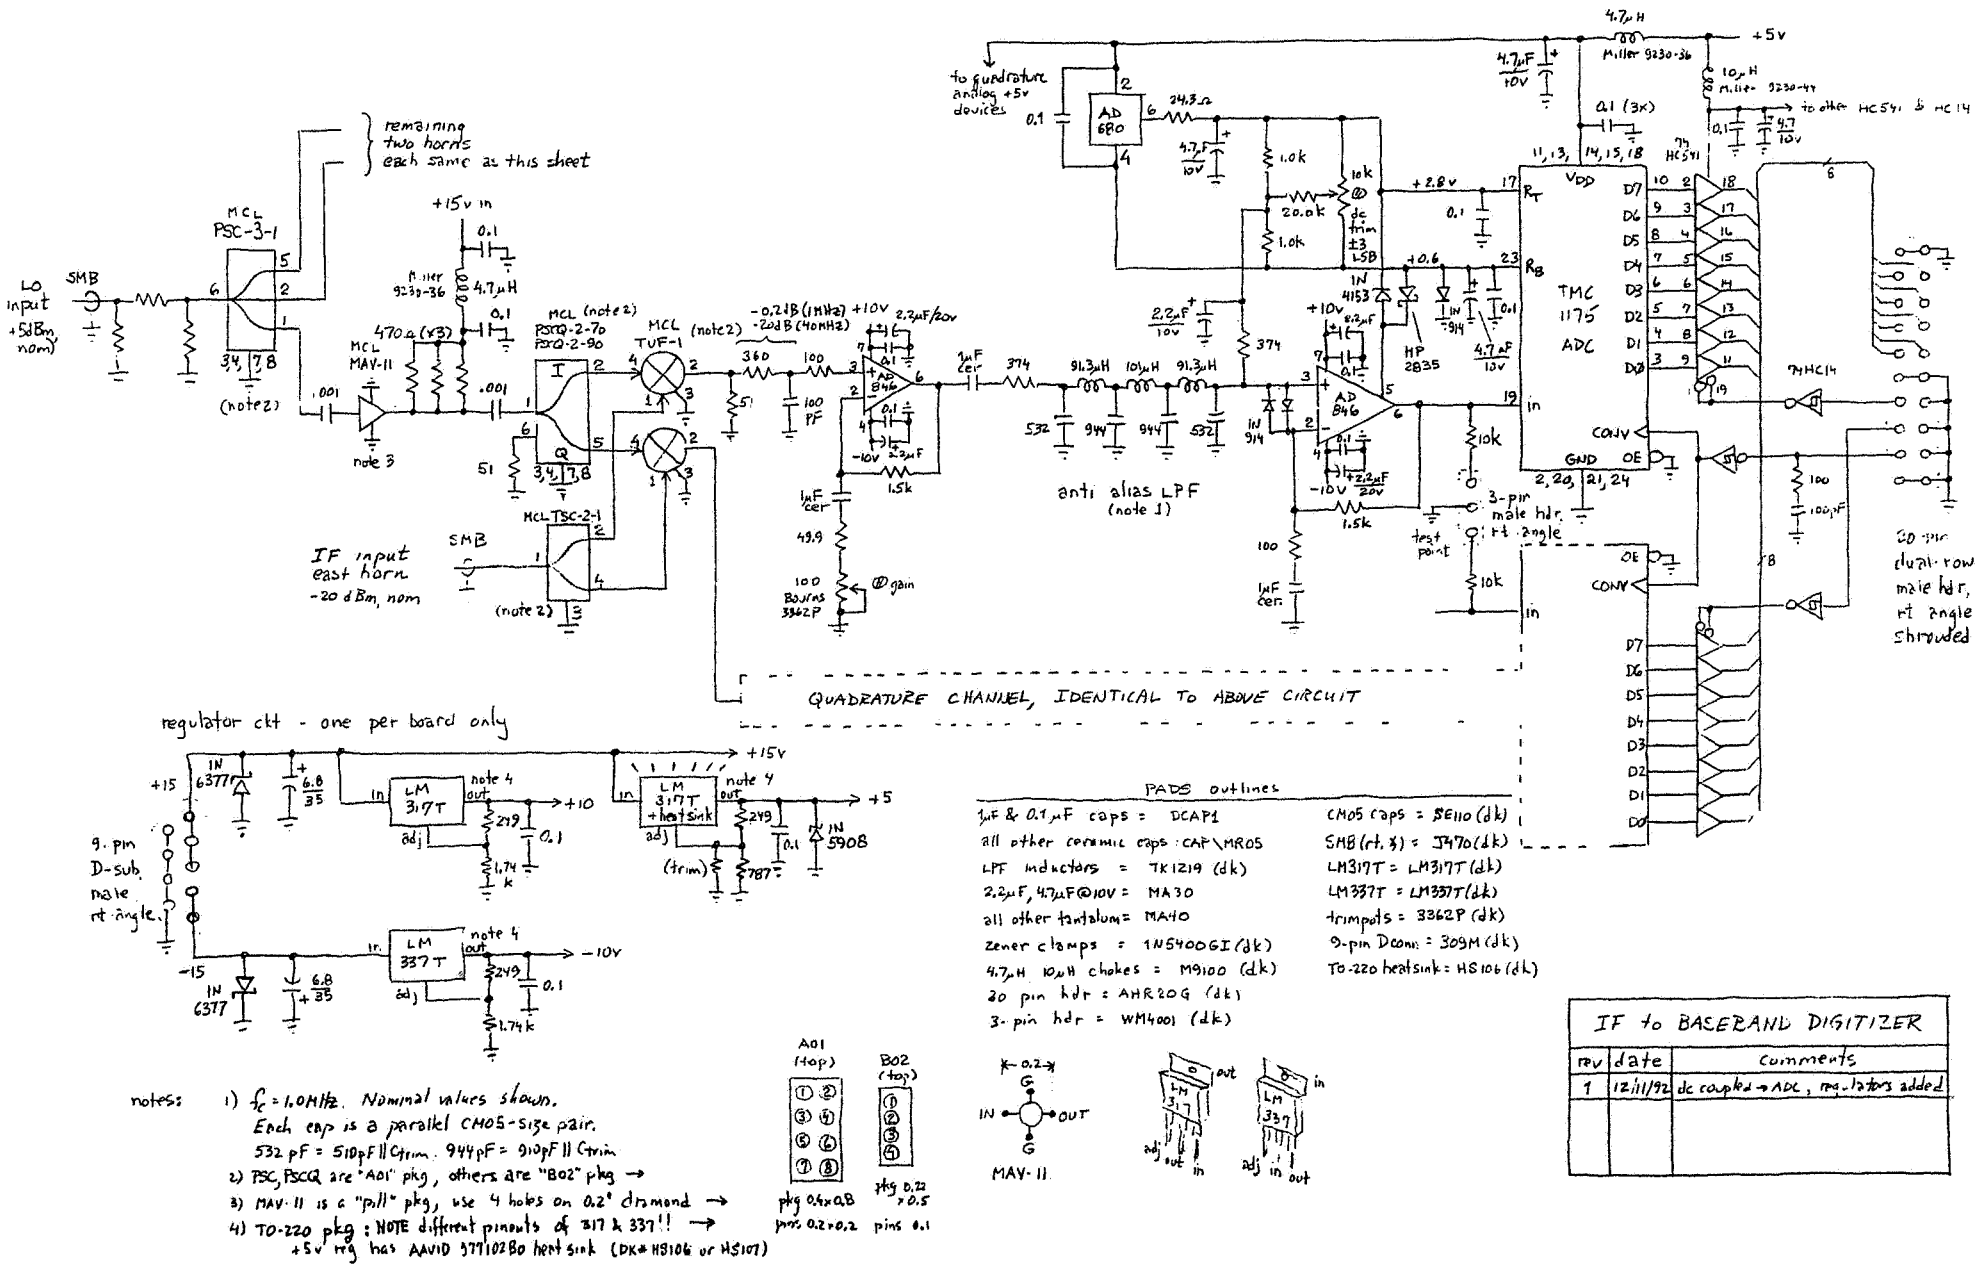

Figure 11. Schematic of 3-horn complex mixer-filter-digitizer board; 20 will be required for the 240-megachannel spectrometer.

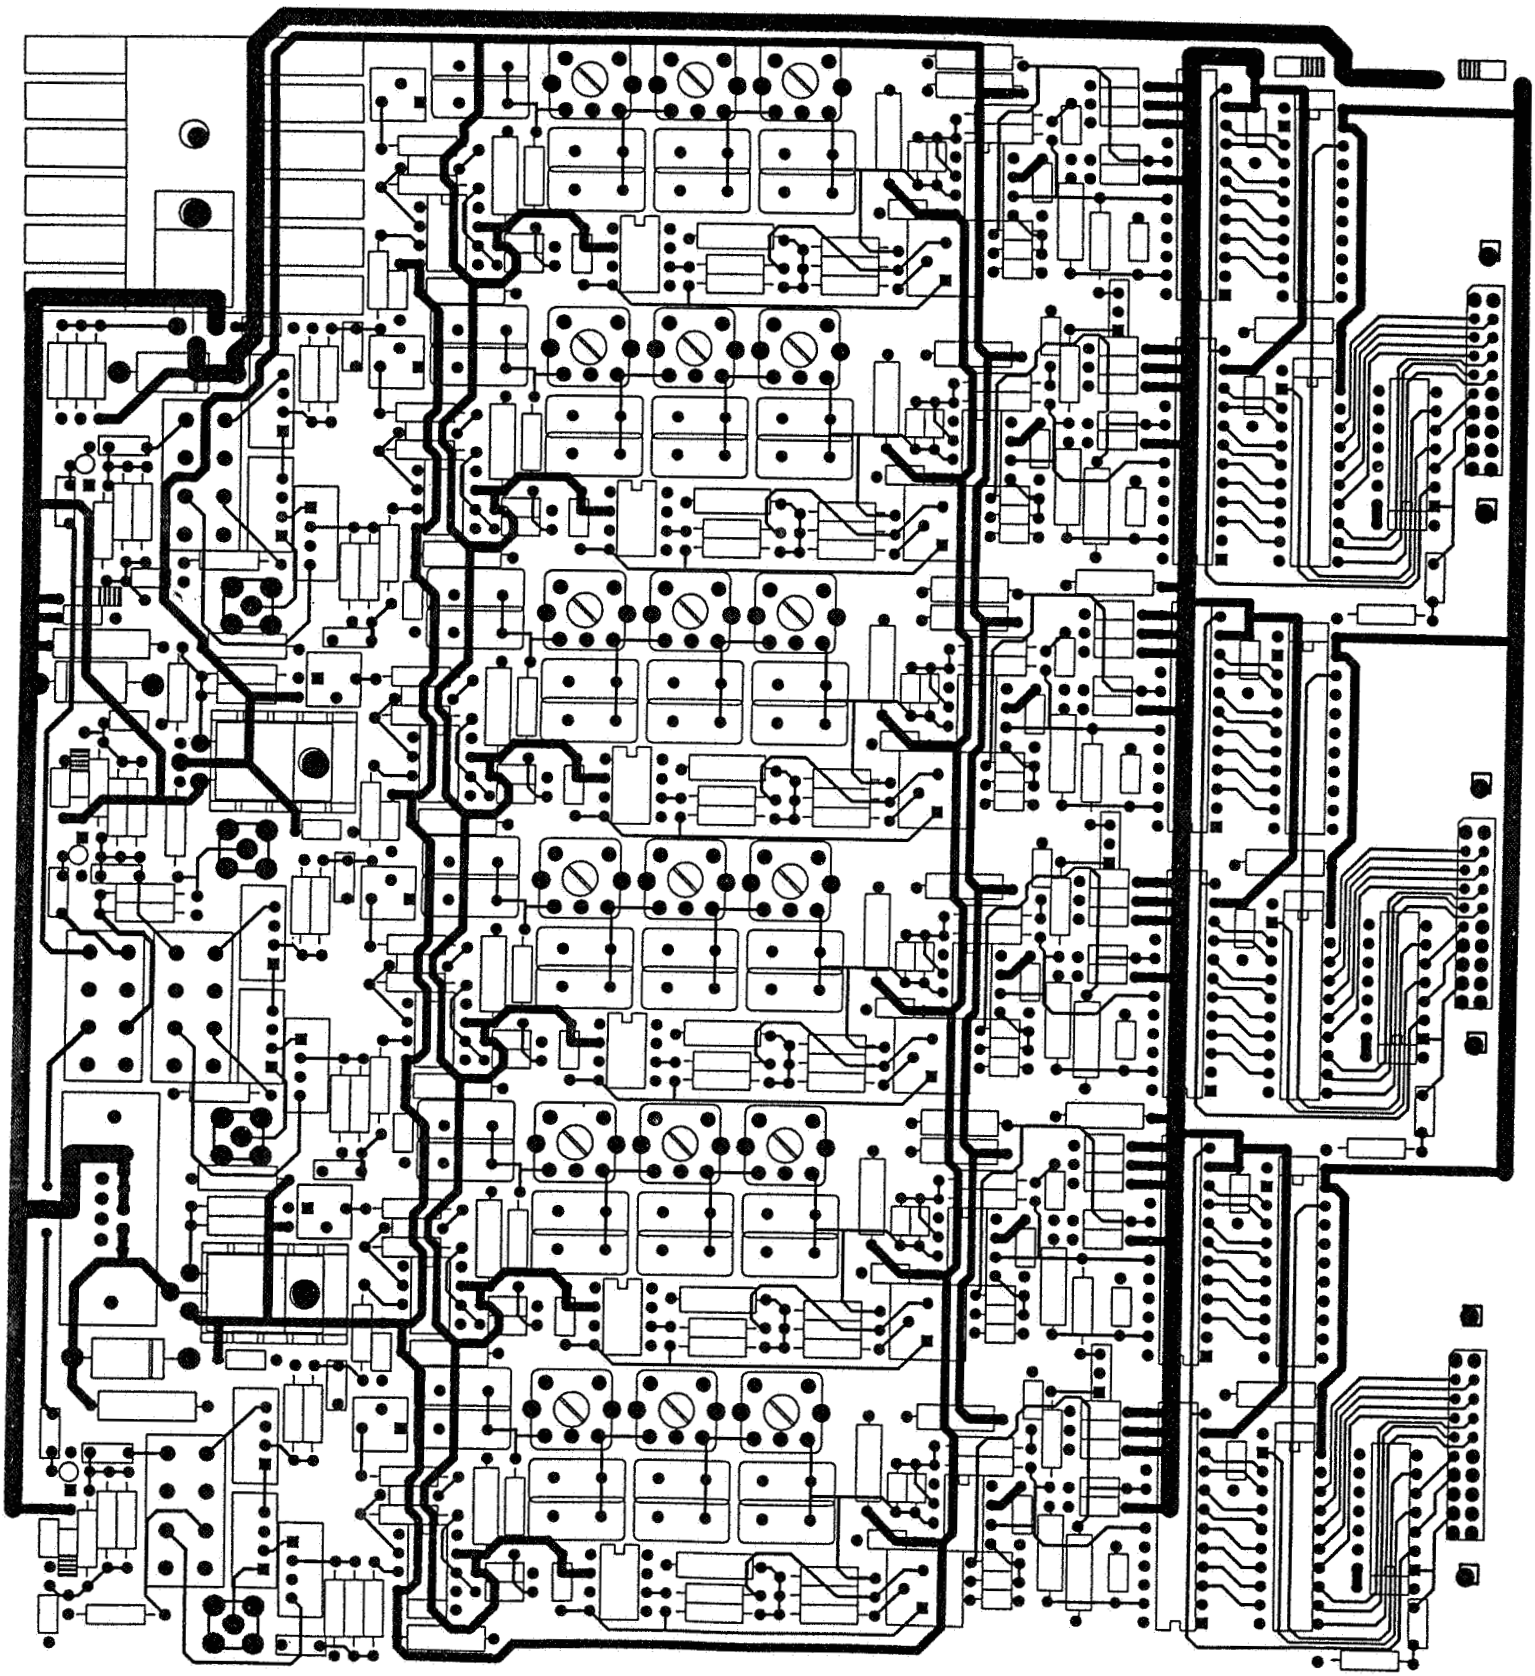

Figure 12. Mixer-filter-digitizer printed-circuit board. Analog signals enter from the left; the anti-alias lowpass filters are at center, followed by baseband amplifiers and A/D converters, finally emerging as 8-bit multiplexed I and Q amplitudes on the 20-pin headers at right. The board is implemented with 12/12 design rules as a 2-sided board with ground plane.

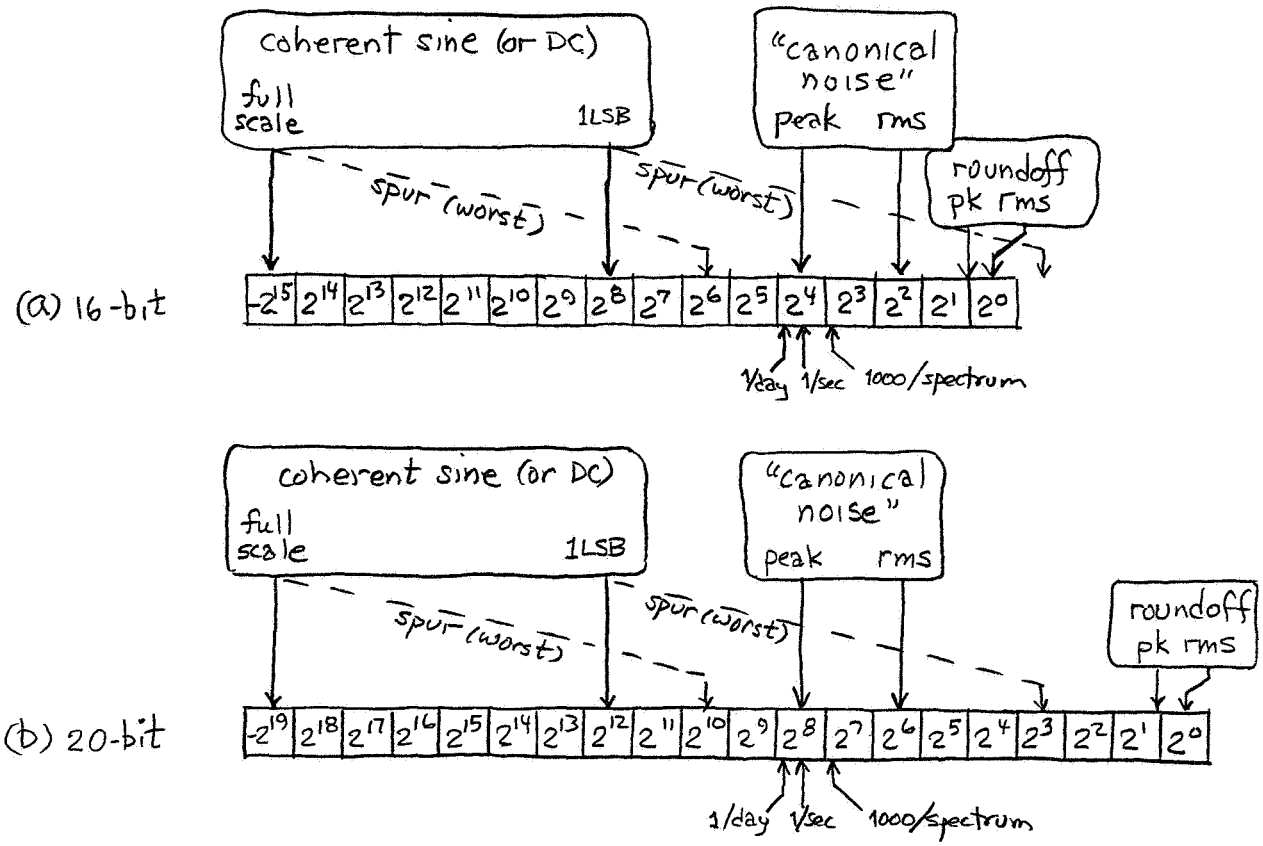

Figure 13. 4-million-point integer FFT behavior with regard to coherent signals, noise, roundoff, and spurs. 8-bit I and Q input amplitudes are assumed, left-justified in the FFT's word, with one righthand bit shift per butterfly, and a 256x256x16 "expanded ROM" for the large twiddle factor. (a) 16-bit integer arithmetic; (b) 20-bit integer arithmetic. The diagrams show the location in the output spectral amplitudes that the indicated inputs emerge. For example, a full-scale input sine wave, whose period is commensurate with the transform window, produces a full-scale output peak in the corresponding frequency bin. "Canonical noise" is input ("antenna noise") Gaussian white noise of amplitude such that approximately one sample in 4 million would overflow full scale (and is forced to saturate at full scale); its modulus has a mode of approximately 20% of full scale. "Spurs" are spurious spectral responses to genuine sinusoidal components in the input time series, caused primarily by the twiddle ROM truncation; each doubling of ROM size reduces them by 6dB; our design uses ROMs twice as large as assumed here, hence produces worst-case spurs that are shifted one bit to the right of the positions shown. Omission of bit shifts between butterflies affects all signals and spurs linearly; however, the effect on roundoff noise is different -- although the *peak* roundoff noise grows linearly with omission of bit shifts, the *rms* roundoff noise grows only as the square root. The three arrows below the boxes point to the thresholds that produce the indicated "hit" rate, assuming random noise, and a 4-million point FFT every 2 seconds.

b3lamp Tue Sep 1 22:29:29 1992 1

100000	-41	8	-9	-6	8	25	-23	-33
200000	5	42	-9	2	7	-37	-29	-10
300000	-17	-25	21	-11	12	15	-31	42
400000	-30	-8	12	-17	12	31	-11	14
500000	-10	13	-11	3	-24	57	23	52
600000	17	-19	1	36	-24	-2	-47	-4
700000	5	3	9	1	-19	-12	-7	4
800000	14	-17	-28	-7	46	-23	-4	-21
900000	1	16	-20	11	14	20	1	-35
999992	503	31	27	13	10	-3	-11	3
999996	11	-28	-14	17	-40	18	19	4
1000000	51422	-20	-4	-25	18	28	-17	-11
1000004	-23	25	5	44	4	9	3	-11
1000008	-14	-81	-1	24	3	-35	-17	-23
1016372	-15	-4	15	-23	8	16	35	-8
1016376	-18	3	3	-21	-35	38	29	2
1016380	10	5	41	22	9	-3	0	19
1016384	-12	1	26	27	-18	10	19	46
1016388	1	1	21	-55	-38	15	21	12
1016392	15	-9	41	23	5	-16	-31	7
1016396	-5	14	11	18	16	-53	-9	16
1100000	-36	7	3	-49	6	-7	7	10
1200000	-22	43	-37	-15	30	-22	-4	-9
1300000	12	1	-12	-43	35	17	-15	37
1400000	3	1	-40	-21	29	11	-19	-15
1500000	-5	-26	-9	-32	0	-31	-5	-17
1600000	17	29	17	60	13	-1	3	-19
1700000	-4	55	47	-23	-5	6	-13	-33
1800000	9	-37	-21	35	-11	-49	21	15
1900000	37	1	30	4	-3	38	20	21
1999972	270	277	283	286	330	305	323	355
1999976	333	121	311	395	370	337	394	419
1999980	423	425	451	439	441	451	469	489
1999984	509	540	507	525	552	633	615	613
1999988	700	733	741	783	816	791	932	922
1999992	979	986	1107	1103	1319	1284	1538	1584
1999996	1871	1920	2531	2574	3607	3663	6543	6619
2000000	32825	32779	10941	10937	-4646	4711	-2955	-2975
2000004	2170	2165	-1743	-1697	-1433	-1434	-1173	-1218
2000008	1040	1075	-903	-920	-817	-899	-711	-764
2000012	731	631	-639	-619	-582	-555	-486	-545
2000016	-538	-487	-475	-459	-443	413	-348	-347
2000020	-107	-53	-525	-493	-421	427	-373	-357
2000024	-355	391	-325	-367	-348	329	-298	-329
2000028	-312	-289	-241	-259	-242	251	-294	-275
2000032	-295	-259	-259	-282	-293	229	-255	-209
2000036	-231	-227	-218	-257	-189	-211	-187	191
2016372	28	-25	33	48	-24	-8	23	11
2016376	-11	-18	41	-2	7	-29	43	-37
2016380	0	1	-24	-45	-6	47	-23	-10
2016384	-26	3	1	58	3	-1	-24	35
2016388	9	31	-23	-82	-17	-25	-19	29
2016392	7	-34	-25	48	19	37	-8	15
2016396	-4	-13	3	-12	-2	-11	12	-15
2100000	8	-33	-24	-8	-23	33	37	15
2200000	21	3	-23	-18	-39	13	12	29
2300000	19	41	41	-24	-25	21	-29	19
2400000	12	-16	2	-16	27	3	-5	-19
2500000	-82	45	-83	-37	-26	21	33	-19
2600000	-5	7	2	13	15	-26	24	11
2700000	9	-25	-32	-23	-48	-25	-17	20
2800000	-71	-19	2	21	6	10	38	-33
2900000	5	9	33	2	-9	8	27	17
2999972	7	609	-25	593	34	633	-9	672
2999976	-12	637	8	675	-16	747	-11	722
2999980	10	836	-17	848	39	905	6	959
2999984	-31	984	-2	1055	-9	1093	17	1226
2999988	37	1332	12	1434	-7	1557	-9	1762
2999992	-20	1950	2	2165	-30	2504	-19	2973
2999996	40	3673	16	4699	-13	6572	-63	10921
3000000	-3	32724	3	-32754	-11	-10903	-32	-6533
3000004	22	-4647	25	-3650	15	-2967	8	-2476
3000008	7	-2179	-18	-1931	7	-1732	-18	-1556
3000012	19	-1426	-12	-1304	3	-1168	37	-1073
3000016	-9	-1014	-31	-900	19	-879	-7	-784
3000020	17	-508	-18	-1141	-9	-879	11	-826
3000024	-12	-750	10	-686	19	-629	7	-633
3000028	-9	-627	13	-667	-12	-558	41	-550
3000032	22	-527	-18	-518	-13	-515	-38	-452
3000036	-3	-447	23	-467	3	-474	-5	-457
3016372	25	-3	22	29	1	11	17	-22
3016376	31	10	-9	27	-74	-29	45	-9
3016380	2	-3	15	-13	29	-19	26	-21
3016384	28	26	21	43	-11	9	-37	11
3016388	-67	1	-8	-71	23	-7	25	-5
3016392	39	8	-18	-9	23	-55	8	49
3016396	-6	23	13	1	-15	-9	-5	13
3100000	19	11	19	5	20	-27	15	-5
3200000	-8	-35	-11	-24	-17	-40	-7	-13
3300000	-4	-11	11	73	18	-3	28	37
3400000	-1	-42	-1	-14	1	-34	11	-23
3500000	11	23	31	27	35	-33	19	7
3600000	11	5	33	13	45	28	78	-13
3700000	19	9	6	-3	-31	42	39	22
3800000	9	21	-37	-27	-5	-6	10	19
3900000	-11	-1	-47	-7	-61	-7	-1	7
3999992	-29	-5	48	21	-27	31	-31	42
3999996	-15	-27	-1	68	-22	0	-37	19
4000000	-111	57	138	-197	39	-37	-15	4
4000004	6	-27	-3	21	1	19	-53	11
4000008	3	7	5	-10	30	16	3	30
4016372	-6	17	52	32	18	11	-2	-5
4016376	-66	19	-9	-1	10	-17	0	-6
4016380	-7	42	37	-2	12	23	14	15
4016384	18	-41	37	1	7	17	8	25
4016388	-4	-38	-12	27	-34	23	-3	39
4016392	15	27	1	33	16	-21	24	6
4016396	5	13	26	6	23	19	3	-21
4100000	33	2	-23	-20	35	17	43	17

window: Uniform, 16K x 16, expanded to 4M

data: noise (0.5 pk amp) + 3 sines (each 0.1 amp) @ 1M, 2M+0.25, 3M+0.5
+ 3 nearby weak sines (each 0.001 amp) @ 1M-8, 2M+20+0.25, 3M+20+0.5
the sum quantized to 8 bits (i.e. each weak sine is 1/8 LSB amplitude)

fft: 20-bit, all scale

twiddle ROM: 128K x 16 expanded to 4M

plot complex amplitude

Figure 14. Result of 4M complex integer FFT, with three strong sines and three weak sines (both "on-bin" and "off-bin") embedded in strong noise (parameters listed on figure), using no windowing. The on-bin strong sine (at 1M) is cleanly resolved, permitting detection of the nearby weak signal (at 1M-8); but the off-bin sines (at 2M+0.25 and 3M+0.5) are unacceptably broadened, by spectral "leakage," burying the weak sines located 20 channels above.

b3iamp Fri Sep 11 17:34:26 1992 1

100000	-18	3	5	-11	11	17	-19	-11
200000	1	4	-7	1	13	-5	-17	-5
300000	-5	3	6	-2	5	-6	-9	23
400000	-15	5	9	-11	2	11	-3	0
500000	1	17	5	-15	-15	9	19	6
600000	5	-21	-1	19	7	-9	-23	-3
700000	11	-8	4	5	-7	-5	0	1
800000	29	-2	-27	5	24	0	-9	-12
900000	1	4	-9	-7	7	11	1	-5
999992	182	-9	-117	13	35	-9	-11	11
999996	9	-14	-271	8	3497	19	12562	-76
1000000	18691	2	12583	-12	3522	19	-279	-15
1000004	-3	5	5	7	-1	-9	1	17
1000008	4	-35	1	35	0	-20	4	6
1016372	7	5	-1	-13	-4	13	15	-11
1016376	-18	13	17	-19	-21	19	19	-10
1016380	-16	1	13	8	15	-11	30	11
1016384	-11	-11	18	11	-17	-15	11	12
1016388	-3	3	4	-19	-13	15	15	1
1016392	-11	-11	9	14	0	-9	-9	1
1016396	5	2	-1	17	3	-30	2	32
1100000	-17	19	10	-17	-3	8	1	1
1200000	9	-9	-17	9	21	-6	-16	2
1300000	15	5	19	-15	23	5	-25	17
1400000	0	0	-6	-9	4	12	-3	-8
1500000	-1	11	-2	3	3	-3	-6	5
1600000	-7	-6	2	14	1	-8	7	-5
1700000	13	31	-16	-27	13	18	7	-7
1800000	7	-27	-5	36	-4	-33	10	21
1900000	6	11	3	-10	13	7	19	2
1999972	3	11	-8	-1	11	-8	-8	19
1999976	11	-27	-14	32	9	-30	-2	19
1999980	-1	-6	7	2	-6	-2	1	-5
1999984	9	17	-7	-27	3	32	-5	-29
1999988	9	15	-4	8	-7	-21	15	18
1999992	-9	-7	-6	2	9	-11	3	19
1999996	-19	-22	-49	-52	1498	1517	7038	-7117
2000000	12934	12913	10631	10603	3777	3755	-441	431
2000004	10	1	-4	8	-7	-9	13	9
2000008	12	14	9	25	-13	-23	23	0
2000012	24	17	14	-17	-17	13	28	-12
2000016	26	7	14	4	1	15	-55	-72
2000020	120	133	-107	-106	39	26	-1	13
2000024	7	12	13	1	17	8	17	-8
2000028	19	3	11	1	5	5	-7	-9
2000032	-5	11	14	-13	14	6	5	1
2000036	2	4	3	-11	1	8	9	-5
2016372	-8	23	15	27	-21	-20	21	14
2016376	21	11	20	10	19	-2	15	-9
2016380	4	21	-6	-30	3	25	15	5
2016384	7	-13	3	23	5	-21	-13	4
2016388	13	20	-9	-27	5	3	-8	23
2016392	14	-31	-17	21	16	-6	-8	2
2016396	1	3	3	2	-5	-1	7	4
2100000	4	7	-1	-3	-7	10	3	-11
2200000	9	1	-4	-6	-11	2	15	5
2300000	2	19	-7	-21	9	7	-14	13
2400000	5	-1	-5	-3	10	5	-9	6
2500000	10	23	-2	-29	-2	25	10	-18
2600000	7	-5	-3	13	-1	-17	5	11
2700000	-3	1	-1	4	-2	-10	-9	15
2800000	-27	-17	19	13	-12	-3	7	7
2900000	17	10	17	-1	19	-2	19	4
2999972	3	17	-18	-11	20	-1	-9	13
2999976	-3	-10	8	-6	-5	22	-3	-29
2999980	13	23	-18	-11	16	2	-1	7
2999984	-10	-10	11	13	-7	-15	0	9
2999988	5	0	-3	3	-3	-10	5	9
2999992	-9	2	11	-7	-11	3	9	-7
2999996	-9	9	2	-35	-3	1213	43	7517
3000000	-3	16951	2	-16957	4	7524	-13	-1217
3000004	10	41	-1	-10	-1	-1	-2	11
3000008	7	-7	-11	3	11	-3	-7	3
3000012	1	-1	0	-5	-7	3	11	5
3000016	-3	-15	-11	18	17	1	-15	-75
3000020	13	175	-8	-177	-1	83	8	-21
3000024	-8	3	3	-3	1	8	3	-9
3000028	-9	15	13	-27	-16	21	11	-9
3000032	2	3	-9	3	10	-11	-9	9
3000036	2	1	5	-1	0	-1	-11	3
3016372	11	-12	1	9	-5	2	0	-9
3016376	3	1	12	13	-33	-13	34	5
3016380	-17	2	2	-2	2	-2	-2	13
3016384	0	-1	3	6	-4	-2	7	-9
3016388	-13	21	7	-25	3	16	-7	7
3016392	16	2	-20	11	13	-27	-4	25
3016396	0	-8	1	1	1	-5	-6	5
3100000	-6	-5	7	7	-3	-5	5	-13
3200000	-1	-14	3	10	-4	-9	9	5
3300000	-5	-26	3	32	-1	-26	15	14
3400000	-1	-4	-1	10	3	-6	-2	2
3500000	10	1	-4	12	9	-17	-11	7
3600000	10	1	-1	-4	28	11	-41	13
3700000	11	1	6	-11	-19	10	10	2
3800000	1	18	11	-13	9	1	-9	7
3900000	1	-1	1	0	-14	1	19	4
3999992	17	-9	28	-1	-17	13	2	-10
3999996	-1	-15	5	29	-8	-18	29	-16
4000000	-64	61	63	-75	-22	37	-7	2
4000004	9	13	-8	9	15	3	-19	10
4000008	13	9	-8	-13	9	-15	-9	-9
4016372	-7	0	15	4	-10	-1	15	-7
4016376	-20	8	11	1	1	-1	2	9
4016380	-11	18	13	-17	-7	13	2	-1
4016384	-5	-9	9	8	-7	6	4	16
4016388	-3	-24	5	12	-8	11	3	-23
4016392	4	11	-5	9	1	-14	5	7
4016396	-8	1	4	6	1	13	0	-15
4100000	27	2	-23	-9	10	8	6	-3

window: Blackman-Harris, 16K x 16, expanded to 4M

data: noise (0.5 pk amp) + 3 sines (each 0.1 amp) @ 1M, 2M+0.25, 3M+0.5
+ 3 nearby weak sines (each 0.001 amp) @ 1M-8, 2M+20+0.25, 3M+20+0.5
the sum quantized to 8 bits (i.e., each weak sine is 1/8 LSB amplitude)

fft: 20-bit, all scale

twiddle ROM: 128K x 16, expanded to 4M

plot: complex amplitude

Figure 15. Result of 4M complex integer FFT, with three strong sines and three weak sines (both "on-bin" and "off-bin") embedded in strong noise (parameters listed on figure), using a Blackman-Harris window (implemented as a 16Kx16 "expanded ROM"). The on-bin strong sine (at 1M) is now somewhat broadened, but not so much as to obscure the nearby weak signal (at 1M-8); now, however, the off-bin sines (at 2M+0.25 and 3M+0.5) are kept reasonably narrow, permitting clear detection of the weak sines located 20 channels above.

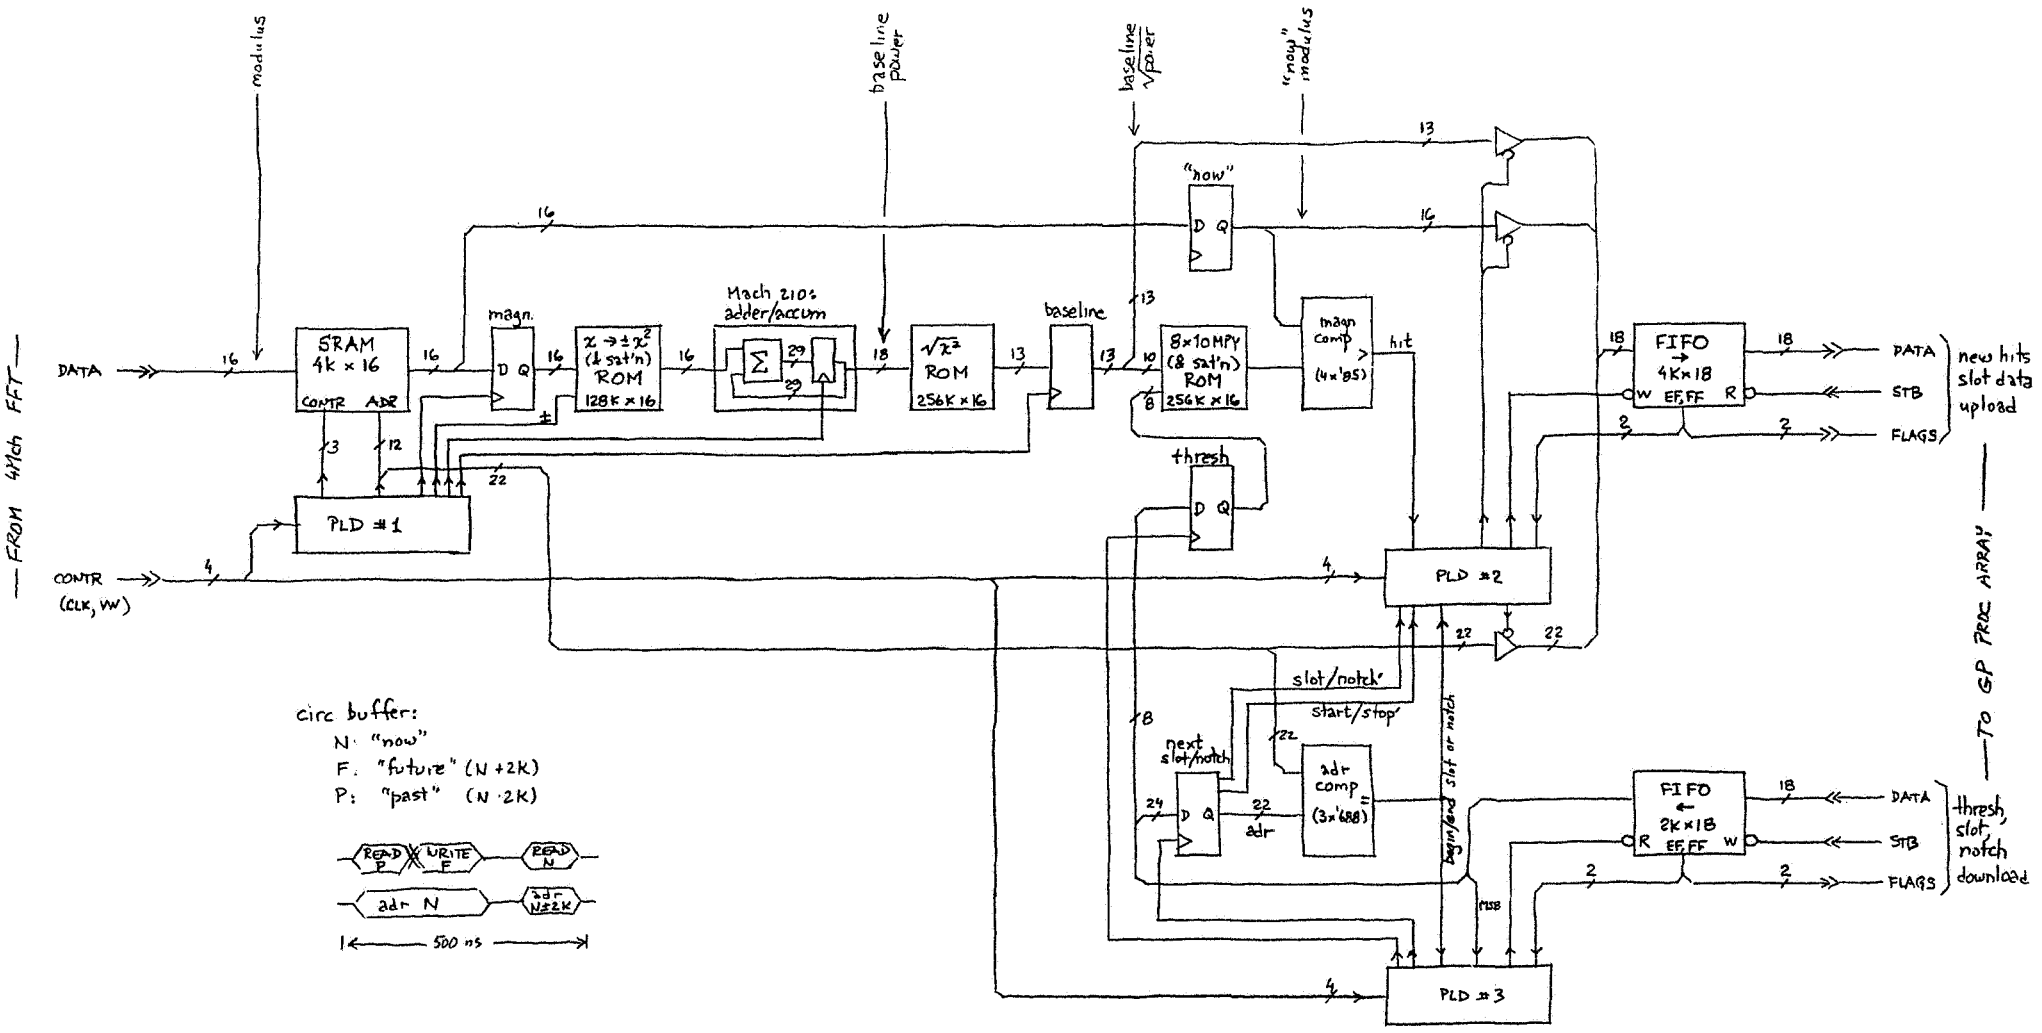

Figure 16. Block diagram of the hardware baseline module. The static RAM ("SRAM") implements a circular buffer, forming a moving average of 4k contiguous data points centered on the middle point ("now"); the accumulator (implemented in a Mach 210 complex PLD) adds the newest point ("future"), and subtracts the oldest ("past"), thus maintaining the moving and centered average. The points being averaged are baseline power, saturated at ≈10dB above average noise to prevent bias by strong spectral features. The baseline, converted to modulus and multiplied by a downloaded threshold multiplier, is compared with each data point (also modulus) to find new "hits", which are reported via the output FIFO. Simultaneously, certain designated frequencies ("slots") are reported, whatever their modulus, and others ("notches") are ignored. The locations of individual slots and notches, along with the global threshold value, are downloaded via the input FIFO. Both slots and notches are specified by beginning and ending frequencies, thus permitting arbitrary widths.