Wide-Bandwidth High-Resolution Search for Extraterrestrial Intelligence

Semiannual Status Report

15 Dec 1992 -- 15 June 1993

Paul Horowitz, Principal Investigator

Harvard University
Cambridge, MA
02138
June 15, 1993

Grant Number NAGW-2872

1. INTRODUCTION

This report summarizes research accomplished during the third 6-month period of the grant. During the period covered by this report the active personnel included the PI, three graduate students (Darren Leigh, Rik Koda, and Jonathan Weintroub, each averaging a quarter-time research commitment), a Harvard undergraduate (Derrick Bass), and a recent mathematics graduate from Harvard (Nick Shectman).

2. RESEARCH ACCOMPLISHED

2.1 Dual-Horn Antenna Performance

As described in the previous progress report, we designed, built, and installed a pair of stacked pyramidal horns. The design was a compromise, with a reduced edge taper of 5.5dB (compared with the canonical 10dB). This was done in order to increase beam lobe overlap, and to permit two horns to fit in the existing radome, which was originally built to hold a single conical horn. We took the unusual step of truncating the corners (see Figure 4 in the previous report), to allow a larger horn aperture within the constraints of a cylindrical radome. Our calculations predicted only a 0.6dB loss of antenna efficiency, and a point-source crossover at the -4dB point; we assumed no feedhorn interactions.

As shown in the previous report, the horns *do* fit (with about an inch to spare!), and the drift scan of an astronomical source exhibits the desired lobe handoff, with crossover at -4.5dB. The signal/noise ratio seemed low, but cold weather made life miserable at the focus, so we did not make further measurements. It was not clear whether the low SNR was due to amplifier and feed losses, or to the compromise feedhorn illumination. In particular, the amplifiers have been at the focus for a decade, use earlier GaAsFET technology, and their noise figure had not been recently measured. Furthermore, the new horns look through the entire radome cover, including an oblique path through the curved edges, which was never intended; the radome material itself is fiberglass, now exposed to the weather for some 20 years. On the other hand, the dual feedhorn design is novel. Thus we did not know the relative contributions of noise temperature and aperture efficiency to the observed SNR.

In the meantime, antenna pioneer Dr. John Kraus became interested in our experiments, and devised a "hot load" technique by which we could unambiguously measure both the system noise temperature T_N (due to amplifier and feed losses), and also the aperture efficiency ε_{ap}. Furthermore, he actually sent us a crate full of "hot load" absorbing material, along with a recipe for its use!

We made the measurements in May, comparing cold sky (with and without the radome cover installed) with the hot load absorber, and with the astronomical radio source Tau A (the Crab nebula). They are summarized in Figure 1. The result is that the antenna efficiency (product of aperture efficiency and spillover efficiency) is as predicted (approximately 45%, compared with the usual 50%), and that the system temperature is a modest 112 Kelvin, of which ≈60K is likely due to the uncooled GaAs amplifiers, perhaps another ≈30K from feedhorn losses (we noted, but could not easily remove, approximately 1/3 "monolayer" of dead flies in the waveguide-to-coax

adapter), and the rest probably coming from antenna sidelobes that illuminate the 300K environment. The radome appears to be entirely benign. This performance is satisfying, and will be improved significantly by the new HEMT amplifiers we have now completed (see next section). In particular, uncooled HEMTs, combined with aggressive fly removal, should trim ≈40K, for a system temperature of about 70K; cooled HEMTs would trim another 20K, producing T_N≈50K.

We consider the antenna system now complete.

2.2 HEMT Low-Noise Amplifiers

We have now completed all four wideband (1.3 GHz to 1.9 GHz) low-noise L-band HEMT (high electron mobility transistor) amplifiers, using kits of parts from Berkshire Technologies (Oakland, CA). Figure 2 shows their measured performance. They achieve 30K or better noise temperature over the waterhole band (1.4-1.7 GHz) when operated at room temperature, and ≈5K when cooled to liquid nitrogen temperature (to which must be added comparable feedline and connector losses associated with the necessary cryostat; thus a realistic cryogenic amplifier noise temperature is ≈10K). Given a realizable ≈40K antenna contribution to the overall system noise temperature, we will achieve ≈70K system temperature with room temperature operation, and ≈50K system temperature with cooled amplifiers, with the current feed system. Given the modest overall improvement one achieves by cooling the 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.

2.3 Downconverter

The downconverter subassemblies are now complete. Recall that 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. The IF feeds to the mixer boards are provided by an "IF Channelizer" module, consisting of IF sub-band filters (40-60 MHz, and 60-80 MHz), splitters, and buffer amplifiers. Short descriptions follow.

2.3.1 Local Oscillator Array

The LO array synthesizes a set of low-noise local oscillator signals, of high spectral purity, at 40, 42, ... 78 MHz, locked to a master 10 MHz oscillator reference. Our implementation uses 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 reports.

The finished LO array is shown in Figure 3, packaged in its rack-mounted enclosure with bar-graph display. We have burned in this unit now for several months, and it works flawlessly.

2.3.2 *Mixer-Filter-Digitizer*

The LO array drives the mixer-filter-digitizer subsystem, implemented on printed circuit boards using discrete mixers, low-noise wideband current-feedback amplifiers, and hand-tuned anti-alias filters (7-pole 0.1dB Chebyshev lowpass). We chose the TMC1175 monolithic 8-bit flash A/D converters, with the AD680 precision reference. Each module accepts a single LO, and processes three independent IF inputs (East horn, West horn, and Terrestrial horn). The design was implemented on a ground-plane printed circuit board, shown in Figure 4.

Figure 5 shows the completed array of 21 boards (20 channels, plus a frequency-agile redundant channel), along with the LO Array and IF Channelizer (see below).

2.3.3 *Digitizer Tester*

We designed and built a simple "digitizer tester," to verify correct performance of the mixer-filter-digitizer. It consists of timing logic to drive the digitizer, parity checking circuitry, and latches and D/A converters to display (on an XY scope) the "un-digitized" waveform (Figure 6). Figure 7 shows a quadrature pair of digitized/undigitized baseband signals processed this way, plotted as *a)* two channels versus time, and *b)* an XY display of real vs imaginary (a Lissajous plot). The accurate quadrature relationship, and lack of artifacts, demonstrates correct operation of the mixer-filter-digitizer. This unit will also be useful in system maintenance.

2.3.4 *IF Channelizer*

The trio of IF signals from the three feeds are broadband (≈100 MHz), and deliberately low-level (≈-20dBm, for good headroom and low intermodulation). The IF Channelizer splits each IF into the 20 individual feeds for the mixers, at the same time filtering into octave bands (40-60 MHz, 60-80 MHz) to preserve headroom and dynamic range in both the amplifiers and mixers. Thus its output consists of 10 signals at -20dBm in each of the two octave bands, for each of the three IF's (see schematic in Figure 8). The channelizer has been implemented using modular amplifiers, splitters, and filters, on a ground-plane printed circuit board. Figure 9 shows the finished channelizer, a gleaming array of gold-plated SMB connectors, interspersed with GPD amplifiers and other modular RF components. It can be seen in its rack packaging, below the mixer array, in Figure 5.

The BETA spectrometer is thus complete from antenna (microwave), through IF (VHF), baseband, and conversion to digital time series. The following sections describe progress in the remaining (digital) portions of the search system.

2.4 FFT Array

As described in previous reports, 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.

We have verified the architecture through numerical simulation. We also carried out exhaustive analysis of the effects of finite word size, both in the FFTs and in the window and "twiddle factor" multiplications. The results are described exhaustively in the 12/92 progress report. The bottom line is that one can use severely truncated window and twiddle ROMs with little effect on dynamic range, and that weak sines are accurately preserved in the presence of strong sines and noise. Our design uses 8-bit input quantization, an 8Kx16-bit "expanded" Hanning (or Blackman-Harris) window ROM, and 20-bit integer arithmetic for the FFT itself, in which a bit scale is performed after each butterfly ("all scales enabled").

The following sections describe progress in the spectrometer.

2.4.1 Spectrometer Breadboard

We had delayed laying out a printed circuit board for the FFT processor until we saw actual A41102 silicon, preferring to concentrate our efforts on the upstream and downstream modules (antenna, downconverter, hardware backend, etc.). Our original order for the FFT chips was placed in December, 1991; they finally arrived in January, 1993!

We decided to begin modestly, with a single-chip implementation of a 256-point, 20-bit FFT breadboard. The rationale was that we needed such a device to test the chips initially, and also for continuing maintenance. Furthermore, it gives us a chance to try out our mixer-filter-digitizer boards. We also wanted to experiment with window ROMs, barrel shifts, and the like. Figure 10 shows the block diagram of the breadboard.

Beginning at the left, the breadboard accepts either the digitized baseband signal (from the mixer-filter-digitizer), or a synthesized set of harmonic "basis vectors" from the 27C010 ROM. In either case the signal can be multiplied by a chosen window function (via DIPswitch) from the selection stored in the 27C32 window ROM. The FFT chip receives its configuration via the 87C751 microcontroller, again selected manually via DIPswitch. The spectral output (either real or imaginary amplitude) is barrel shifted, permuted, and displayed on an XY scope via 10-bit D/A converters. The breadboard performs 7,812.5 spectra per second, plotting 2 million points per second; you don't see any flicker in this display!

In "basis vector" mode, each basis vector (one of the 256 sines that are commensurate with the time window) is presented repetitively for 2048 successive spectra, producing a "walking spike" display. Because the basis vectors are window-commensurate, they fall in the center of each bin. Windowing broadens them slightly.

In external signal mode, a mixer-filter-digitizer board provides the input signal, typically fed from a fixed LO and a tunable VHF oscillator. Figure 11 shows the kind of plot you get, in this case using a 60 MHz LO and a nearby carrier. We deliberately used a rectangular window, and set the oscillator to be midway between bin centers, to demonstrate the sinc^2 spectral "leakage" effect. Turning on any of the windows causes the leakage to disappear on this (linear) display.

By using the barrel shift one can have a closer look at the sidelobes, which behave as theory predicts; in particular, they are suppressed best by the Blackman-Harris window, with the Hanning a close second. The triangular is poorer, and the rectangular window is a disaster. See Figure 12, in which the output amplitude has been barrel shifted to show the low-level artifacts. In these plots the spectral peaks thus overflow the D/A converter (though the spectral arithmetic is without error), wrapping through the display vertically; the signal has been placed in its worst-case location, at mid-bin.

2.4.2 Radioastronomy Spectrometer

Prior to designing the gigachannel spectrometer for SETI we are working on a similar, albeit simpler, instrument with applications to some interesting and high-profile radio astronomy. In collaboration with Bernie Burke of MIT and Mike Davis of NAIC, we are developing a Spectrometer/Power Accumulator based on the Austek FDP. The instrument will be used in a search for primeval hydrogen at redshift z = 5 or higher, using the Arecibo radiotelescope. The work is funded by the NSF (through a grant to Burke/MIT), the Planetary Society, and the Bosack-Kruger foundation. We have submitted a proposal with Burke and Davis to NAIC, and plan to field the experiment in early 1994.

The objective of the experiment is to look for the hydrogen raw material in the early universe before star and galaxy formation. Very little is known about the state of the universe between the recombination era at z = 1000 and the era of galaxy formation. In fact, the era of galaxy formation is not well defined, but there is widespread speculation that it occurred at redshift greater than 4 or perhaps even greater than 5. All searches for condensation conducted to date have yielded unconfirmed or negative results. Compared with these other searches our work utilizes a larger aperture and a better spectrometer.

The early hydrogen spectrometer is an array of 16 FFT boards for each antenna feed, each FFT fed by one of the 2.5MHz baseband channels of the existing filter bank. Each quadrature analog signal is digitized to 8-bit integer precision, then multiplied by a window function and Fourier transformed by the Austek FDP. The FFTs are 256 points in length and are computed at a data rate of 2.5 complex megasamples per second, or 10^4 spectra per second. The resolution binwidth is thus 10kHz. The complex spectra are converted to squared modulus in a 16x16 multiplier/accumulator (MAC), with accumulated spectra stored temporarily in a 256x32 bit fast SRAM buffer.

We have cooperated closely with Jim Cordes of Cornell to design the spectrometer with features making it useful for pulsar studies as well. A flexible command interface permits the integration time to be shortened to as little as a single spectrum (0.1ms), the FFT length to be chosen as any power of 2 from 2 to 256 points per board, the time windows to be interleaved from board to board, and the output to be taken optionally as complex amplitudes instead of power. Compared with the requirements of the early hydrogen experiment, data rates are typically four orders of magnitude higher in dynamic pulsar work, and the output interface has been designed to support these rates. Other applications for the equipment are likely to arise, and it is expected that this spectrometer will become part of the permanent equipment installation at Arecibo.

This spectrometer is a concrete example of technology developed for SETI providing benefits in other areas of science. The design spins off neatly from our SETI efforts and fits in well with our overall objectives. Our participation in the experiment better balances the work of our group by facilitating our participation in leading edge radio astronomy and cosmology. The project provides solid material for a graduate thesis, and the expertise developed by the group in designing this spectrometer is directly applicable to, and a warmup for, the design of the extremely complex FFT core for the SETI apparatus.

We have completed the system design of the instrument, as well as the detailed design of the 256 point Spectrometer/Power Accumulator circuit. Figure 13 is a scaled down copy of the schematic; design of a four layer printed circuit board is progressing well.

2.4.3 Four-Megachannel FFT Spectrometer

The Arecibo 4096-channel wideband spectrometer -- our first complete FFT-based spectrometer -- is a finished schematic diagram and netlist, now being routed for printed-circuit manufacture as a 4-layer board. In parallel with that work we are now converting the BETA spectrometer block diagram (Figure 14) to a full schematic. The basic architecture, as shown in the figure, was described in previous reports, where we mentioned that we have devised an interesting compression scheme for reducing memory requirements in the final data permute. (That permute is necessary in order to restore contiguity of the spectrum, so that baselining and thresholding can be done in a single pass in the feature recognizer.) We would like to describe our somewhat unusual compression.

Figure 14 shows a ROM-based data compression of the spectral amplitudes to 16-bit quantities (both I and Q are 20 bits, thus a lossless final permute would require 4Mx40 bits, or 32 MByte), thus requiring only 40% as much permute memory. We have come up with a quasi-logarithmic compression scheme that may be original, and that appears to be optimal for compression of integer data with uniform worst-case fractional error.

The scheme is diagramed in Figure 15. We begin by accepting saturation ("clipping") of large amplitude signals, on the assumption that, once a signal is very much larger than the average noise level, it is *interesting*, regardless of its exact amplitude. Thus, we begin by clipping the 20-bit signed amplitudes to 16-bit signed integers (the opposite of sign-extension -- a sort-of "sign de-extend" operation). Since we are ultimately computing and storing an approximation to signal *magnitude* (for easy comparison with a threshold that is a multiple of the baseline magnitude), we now strip the sign, and pose the task of converting a complex unsigned amplitude pair, I and Q (each now 15 bits), to a 16-bit magnitude. Of course, we could imagine squaring and adding, then computing the square root; but a ROM-based table lookup is simpler, and much cheaper. The problem is that the address is 30 bits, if we naively use the I and Q magnitudes to form an address.

Hence the need for compression: If we can trim the I and Q magnitudes to 9-bit quantities (not necessarily linear), the combined pair is 18 bits, the address space of an inexpensive 256K ROM. For natural reasons, we would like the compression to incur approximately uniform *fractional* error.

We began by imagining a 9-bit floating point format (method I), with 4 bits of exponent required; this incurs a worst-case error of 3.1%, which however can be reduced to 1.6% by the old trick of a "hidden bit." We quickly realized that, even with the hidden bit, floating point format is not ideal: only 10 of 16 possible exponents are used; furthermore, the fractional error varies by a factor of two, depending on the bit pattern of the "fraction" portion (sometimes incorrectly called the mantissa).

We therefore explored a pure logarithmic compression (method II). After all, what else could produce a uniform fractional error? Logarithmic compression is used in audio analog/digital conversion, for just this property. In our case a logarithmic representation does work better (1% error), but it is not optimal -- in particular, there are "missing codes" at the low end (see Figure 15). One third of the compressed codes are squandered on the first 0.1% of the input domain. (Note that this flaw is absent if the input variable is continuous, as in the audio ADC.)

Our third method fixes this flaw, by using a hybrid linear-log representation: Map small integers to themselves ($p(n)=n$), continuing up to a value n_0 at which the fraction integer step equals the fractional error that would be incurred by switching over to a logarithmic compression of the remaining input values. This eliminates missing codes, and appears to be optimal. For our parameters the crossover occurs at n_0=70, and the worst-case error is 0.7%. The original floating point format (with hidden bit) is only half as good, by comparison. We have used this compression in our feature-recognizer simulations, and, not surprisingly, it works as predicted. We have not seen this method mentioned or used elsewhere, though we find it hard to believe that such an obvious technique is novel.

In terms of implementation, the 4-megachannel spectrometer is the most complex module in the BETA system. We anticipate that this major task will occupy several months at least, with prototype board manufacture late in the year.

2.5 Backend "Feature Recognizer" Array

The backend baseline and peak detector (which we now prefer to call a "feature recognizer") was described in the last progress report. We have decided to implement this array of 63 boards in discrete hardware, rather than with commercial DSP chips, for reasons of cost and performance: a single $250 hardware board can outperform an implementation using a DEC Alpha, the fastest commercially available processor. Of course, dedicated hardware is inflexible, by comparison. The feature recognizer boards are each fed by a single FFT board, and reside in groups of three (corresponding to the three antenna feeds) on PC-compatible motherboards (the "general-purpose array") that communicate via Ethernet with the central workstation.

The design is complete at the block diagram level, and could be implemented at the gate level at any time. We are delaying this commitment, however, because we believe that the design may change as the general-purpose array interface evolves during its development, and we do not want to be locked into an inflexible hardware design too soon.

2.6 General-Purpose Array and Workstation

As described in the previous report, the general-purpose (GP) computer array will accept baseline and hit data from three of the "feature recognizer" modules (a single 2 MHz sub-band from the three feedhorns) and do preliminary analysis on it. The GP array will consist of twenty Intel 80486 based PC motherboards, communicating with the central controlling workstation over a thin-wire ethernet. By putting so much hardware horsepower ahead of the workstation and running special DSP algorithms on it, we should be able to look at signals closer to the noise floor, making the instrument more sensitive.

We have written software libraries to access and control the ethernet hardware on both our Unix workstations and the PC motherboards. We have also done extensive testing with them on our local ethernet and have verified that the speed and reliability are sufficient for GP array communications. Since the motherboards will be diskless, we have been developing software to allow them to boot from the workstation over the network. This will consist of a server application on the workstation, and a special boot ROM for the ethernet cards on the PC motherboards. The workstation server has been completed and the boot ROM code will be completed in the next week or so.

We are currently working on the signal processing software for each of the computers in the GP array. This should be completed by the end of July. Also to be completed this summer is a complete mathematical analysis of the system's response to both coherent signals and noise. This analysis will help us to find optimal algorithms for detecting ETI signals and discriminating them from terrestrial interference.

There has been much progress on the backend workstation. Using a preliminary network and data protocol, we have been running simulations between two workstations, one of which mimics the GP array. These simulate the data flow and contention issues that will occur with twenty PCs communicating with the same workstation over the same network. Preliminary results indicate that there will be no problems. We have also been exploring and modifying higher-level signal detection methods using a genetic algorithm. This algorithm has spent hundreds of hours of CPU time on our Sparcstation trying different detection parameters and fine tuning them. It is our hope that the results of these experiments will help us to optimize the performance of our ETI search.

3. NEXT STEPS

We have made major progress during this 6-month period, on the antenna, amplifiers, downconverter, channelizer, digitizer, FFT breadboard, and radioastronomy spectrometer. The first six are now complete, along with appropriate testing modules. The radioastronomy spectrometer was not part of the grant proposal, but is being pursued (entirely with non-NASA funding) as an irresistible side benefit of our advanced SETI instrumentation. (There is often much talk about the benefits from SETI for radioastronomy; this one is a tangible example.) The impact for BETA SETI is both positive and negative -- it demonstrates technology transfer from the SETI effort, it gives us valuable experience with a simpler spectrometer, and it provides a widely recognized scientific research goal as part of the graduate program; on the other hand,

time spent pursuing radioastronomy is time lost from SETI. We are attempting to balance these factors as we judge best. We believe that the radioastronomy spectrometer will be operational late in the year, around the same time the 4-megachannel spectrometer will be in debugging phase. Our current "pert chart" puts the completion of BETA's instrumentation in late spring, 1994. To aid in achieving that goal, the P.I. is taking a full-year sabbatical "in place," in order to devote full time to SETI. Given our expanded (and still expanding!) goals -- 240 million channels and three antenna feeds, compared with the 100 million originally proposed; and an independent radioastronomy program at Arecibo in search of highly redshifted neutral hydrogen -- we are very pleased with our current rate of progress.

4. OTHER FUNDING

During the period of this report we have received funding from The Planetary Society, the Bosack/Kruger Charitable Foundation, and Dr. John Kraus, in addition to our grant of partial support from NASA.

5. PUBLICATIONS AND TALKS

The Astrophysical Journal has accepted our paper, "Five Years of Project META: An All-Sky Narrowband Search for Extraterrestrial Intelligence," by Horowitz and Sagan, for publication in September.

The Principal Investigator chaired a session at the Seattle meeting of the American Physical Society on "Instrumentation for SETI," at which he gave a talk entitled "Project BETA: A 240-Million-Channel Wideband Spectrometer and Backend." He gave talks at the Department of Electrical Engineering at the University of Victoria, and at Lotus Development Corporation in Cambridge, MA. He also wrote a short piece that featured student participation in SETI for Odyssey magazine, a publication for school children.

We continue to enjoy media coverage, variously in newspapers, magazines, radio, and television. We've stopped keeping track of all this activity, we just do our part by answering their questions!

6. ACKNOWLEDGEMENTS

We are grateful to Bernie Burke (MIT), Jim Cordes (NAIC Cornell), Michael Davis (Arecibo Observatory), Jon Hagen (NAIC), John Kraus (Ohio State University), Dan Werthimer (UC Berkeley), Bill Yerazunis (DEC), and George Zimmerman (JPL) for valuable discussions.

Antenna Calibrations 5-31-93

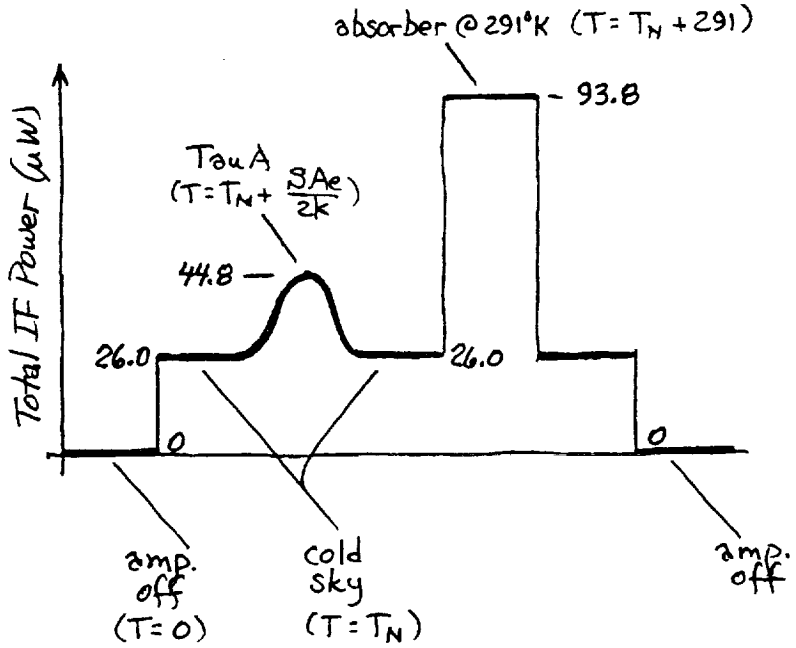

I. System noise temperature, T_N :

$$\frac{T_N}{26.0} = \frac{T_N + 291°K}{93.8} \quad \Rightarrow \quad T_N = 112°K$$

II. Antenna effective area, Ae :

$$\frac{T_N}{26.0} = \frac{T_N + \frac{S_{TauA} \cdot Ae}{2k}}{44.8}$$

$k = 1.38 \times 10^{-23}$ joule/K

S_{TauA} (@21 cm) = 939 Jy

$$Ae = 0.72\, T_N \cdot \frac{2k}{S_{TauA}} = 237\ m^2$$

antenna efficiency $\epsilon = \frac{Ae}{\pi R^2} = 0.45$

III. Sensitivity (K/Jy) $= \frac{Ae \cdot S(1\,Jy)}{2k} = 86\ mK/Jy$

Figure 1. Calibration of system noise temperature (T_N) and effective area (A_e) of each feed of the new dual-feed 26-meter Cassegrain. Bolometric total-power measurements were taken at IF, using cold sky, Tau A, and absorber material at ambient temperature. The Tau A flux value is taken from Baars et al., *Astron. Astrophys.*, **61**, 99 (1977).

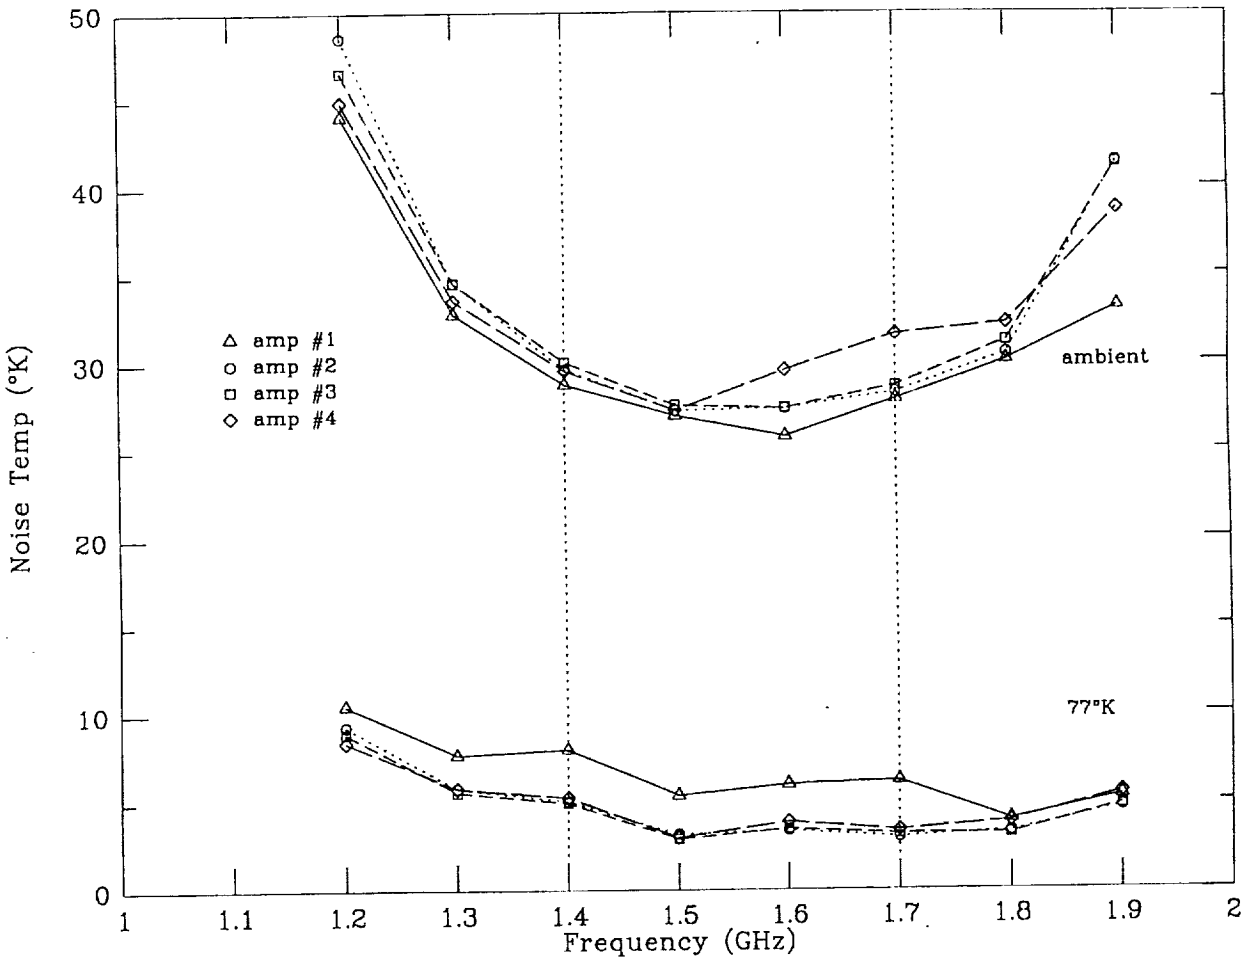

Figure 2. Noise temperature vs frequency for the 4 completed L-band HEMT amplifiers, at 77K and at room temperature. "Waterhole" performance is approximately 5K and 30K, respectively.

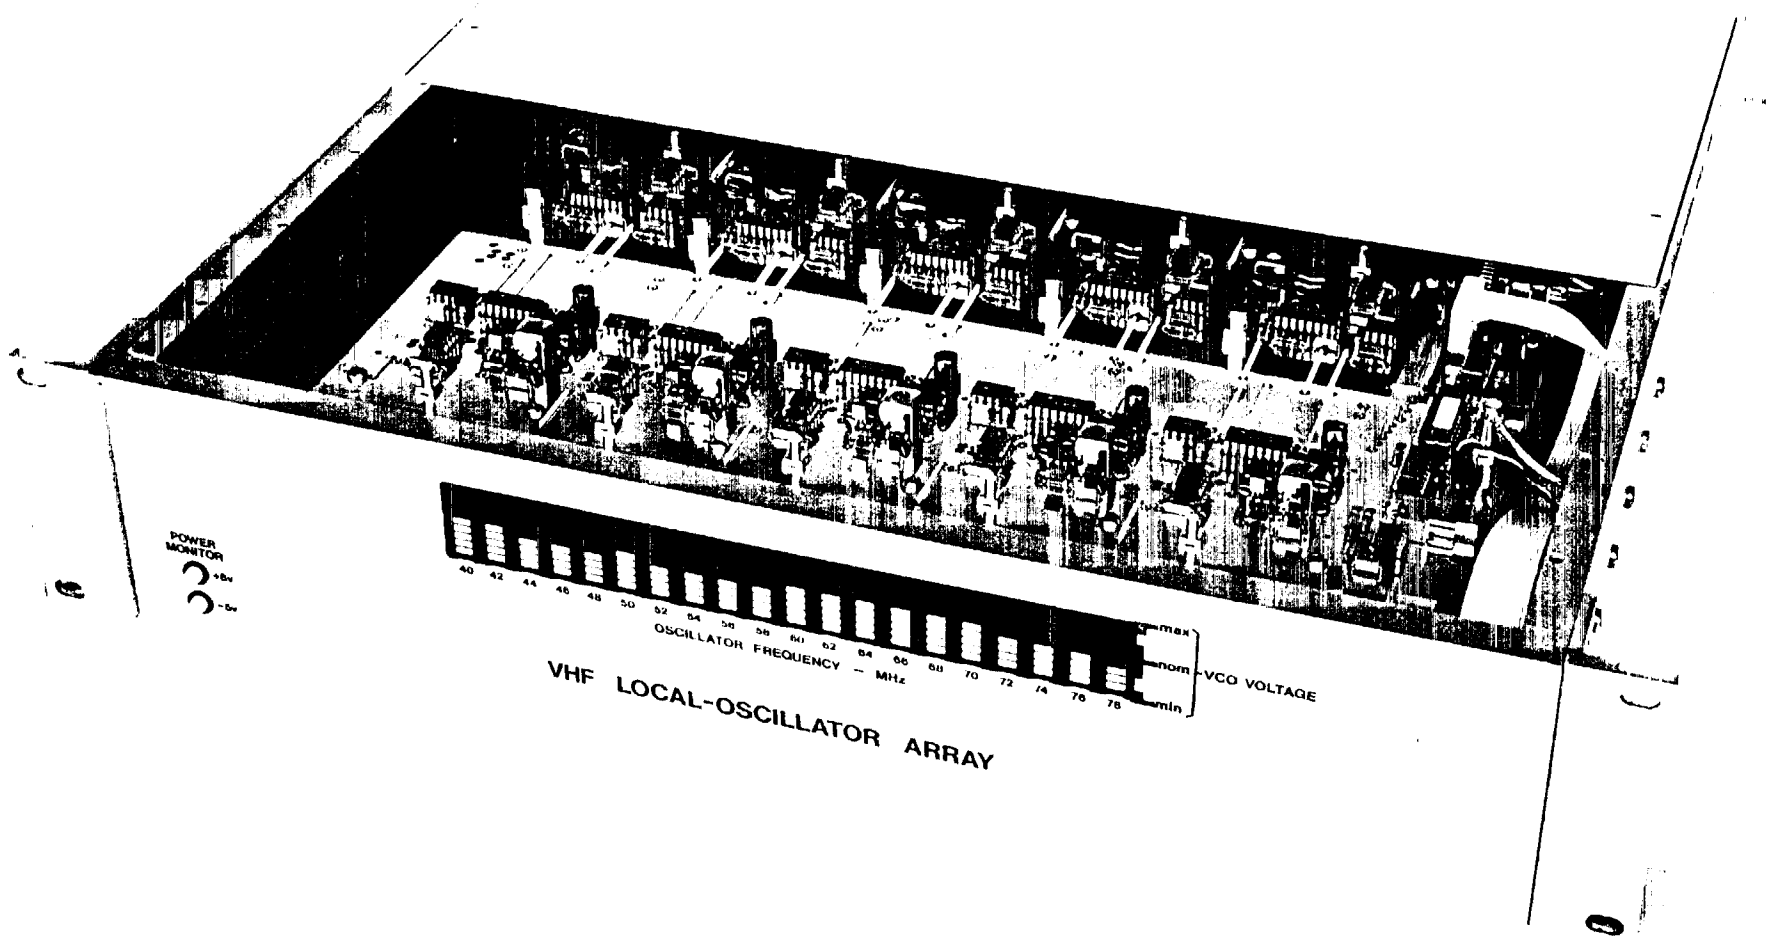

Figure 3. Phase-locked high-purity LO array, packaged in its rack enclosure with front-panel VCO bar-graph display. A second board on the underside completes the array of 20 oscillators.

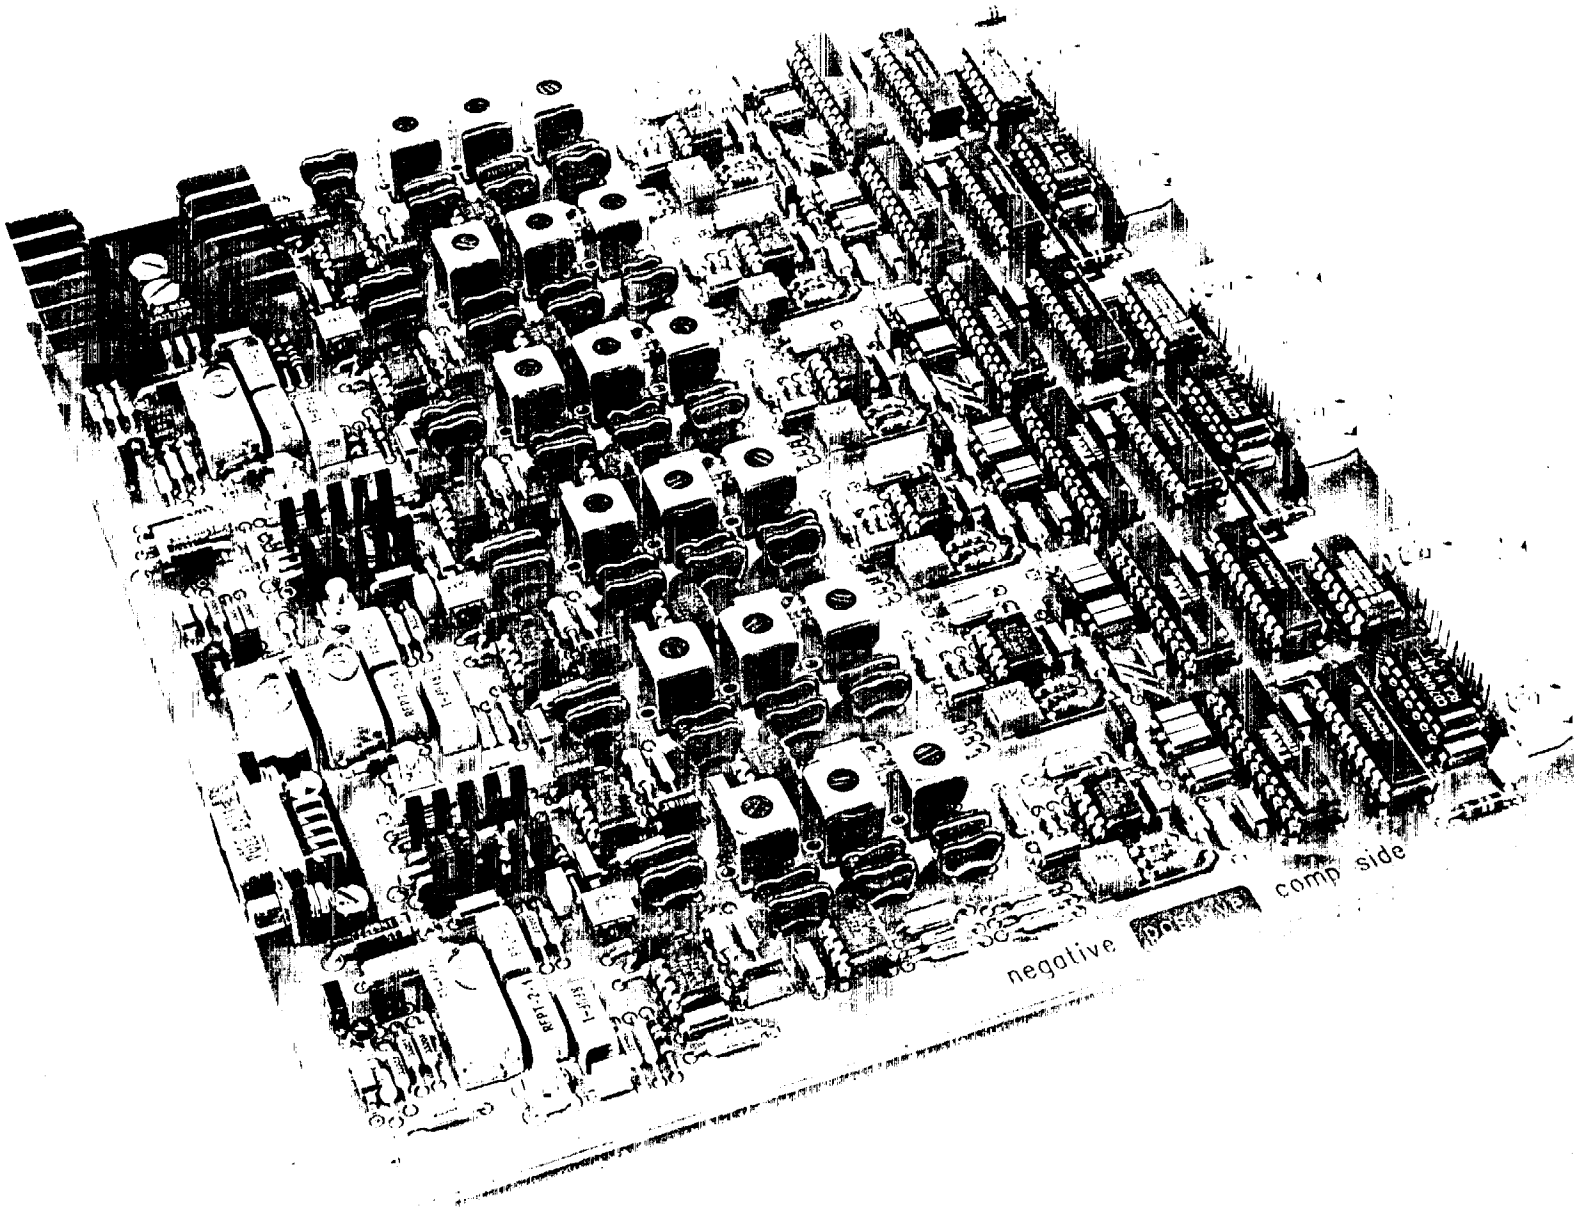

Figure 4. Mixer-filter-digitizer board. A single LO processes 3 independent antenna feeds (East, West, Terrestrial), producing 3 digitized time series. IF and LO signals enter via SMB connectors at left, are mixed, amplified, anti-alias filtered (in the prominent slug-tuned filter at center), conditioned, and digitized in 8-bit flash ADCs, finally emerging via 20-line ribbon connectors at right.

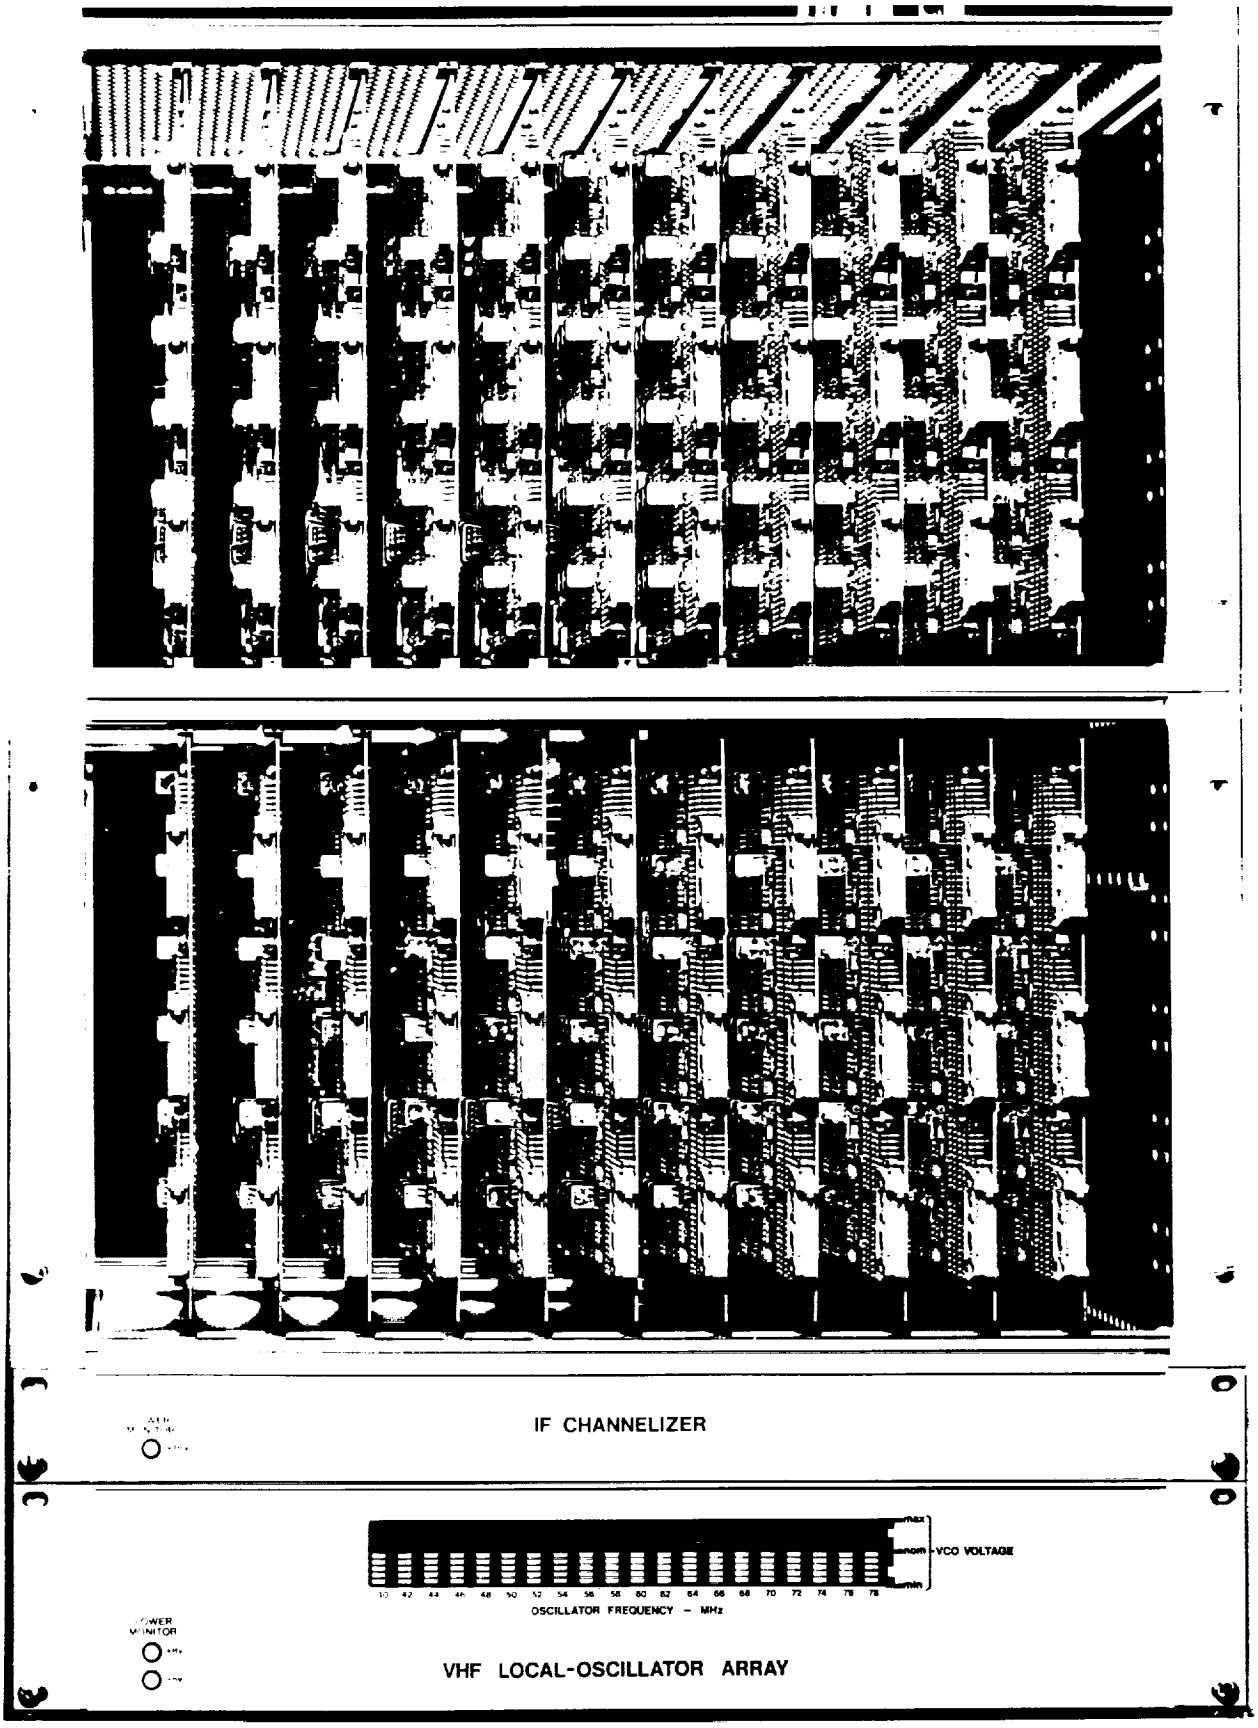

Figure 5. Array of 22 mixer-filter-digitizers in its rack enclosure, along with associated LO Array and IF Channelizer.

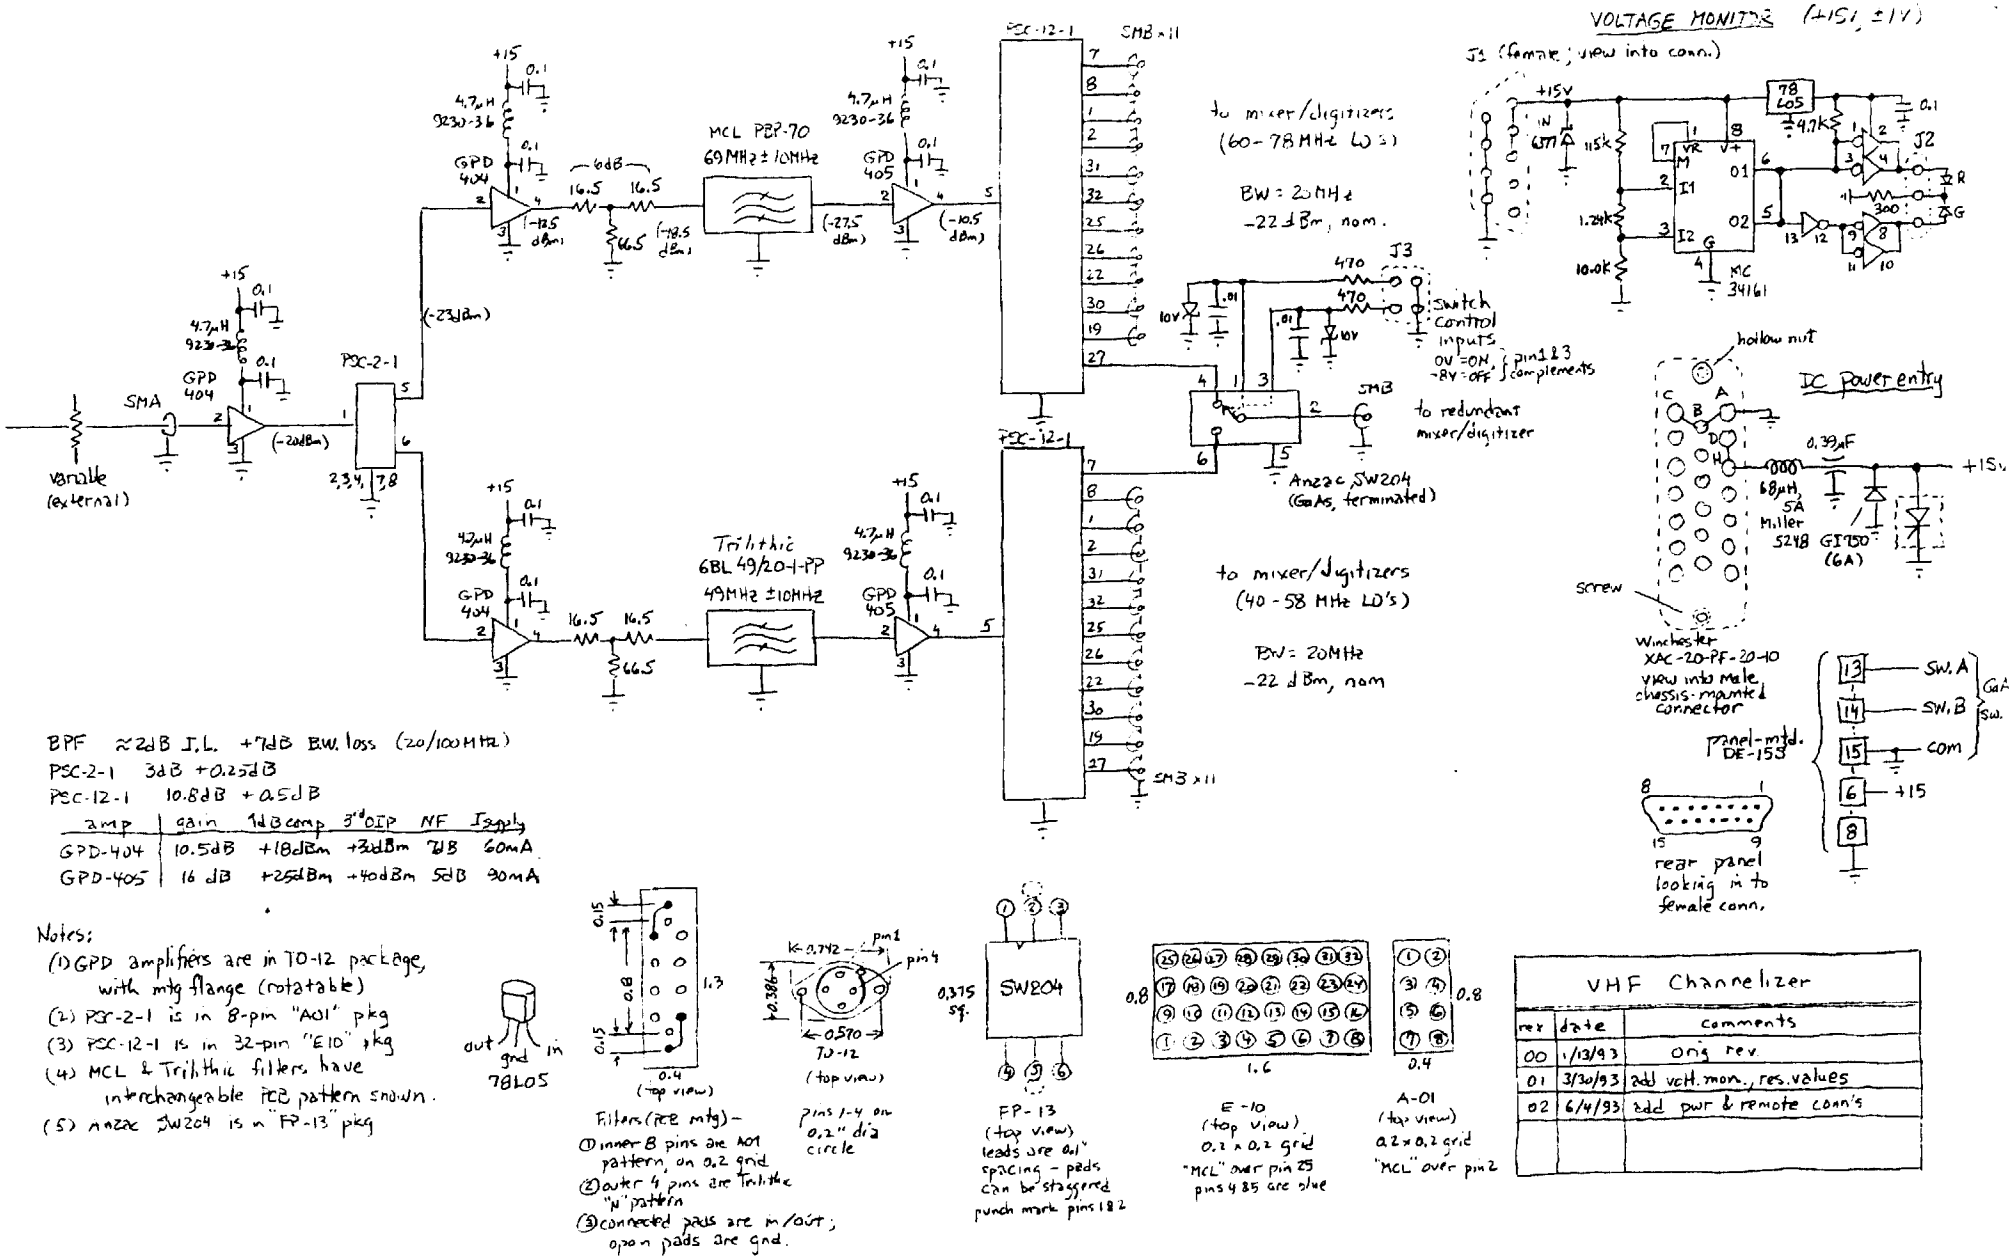

Figure 6. IF Channelizer schematic, showing splitters, octave bandpass filters, and high dynamic range buffers. The outputs feed the mixer-filter-digitizers, with the GaAs switch providing an additional band-switchable feed for the system's redundant Fourier channel. One IF channel of the three antenna feeds is shown.

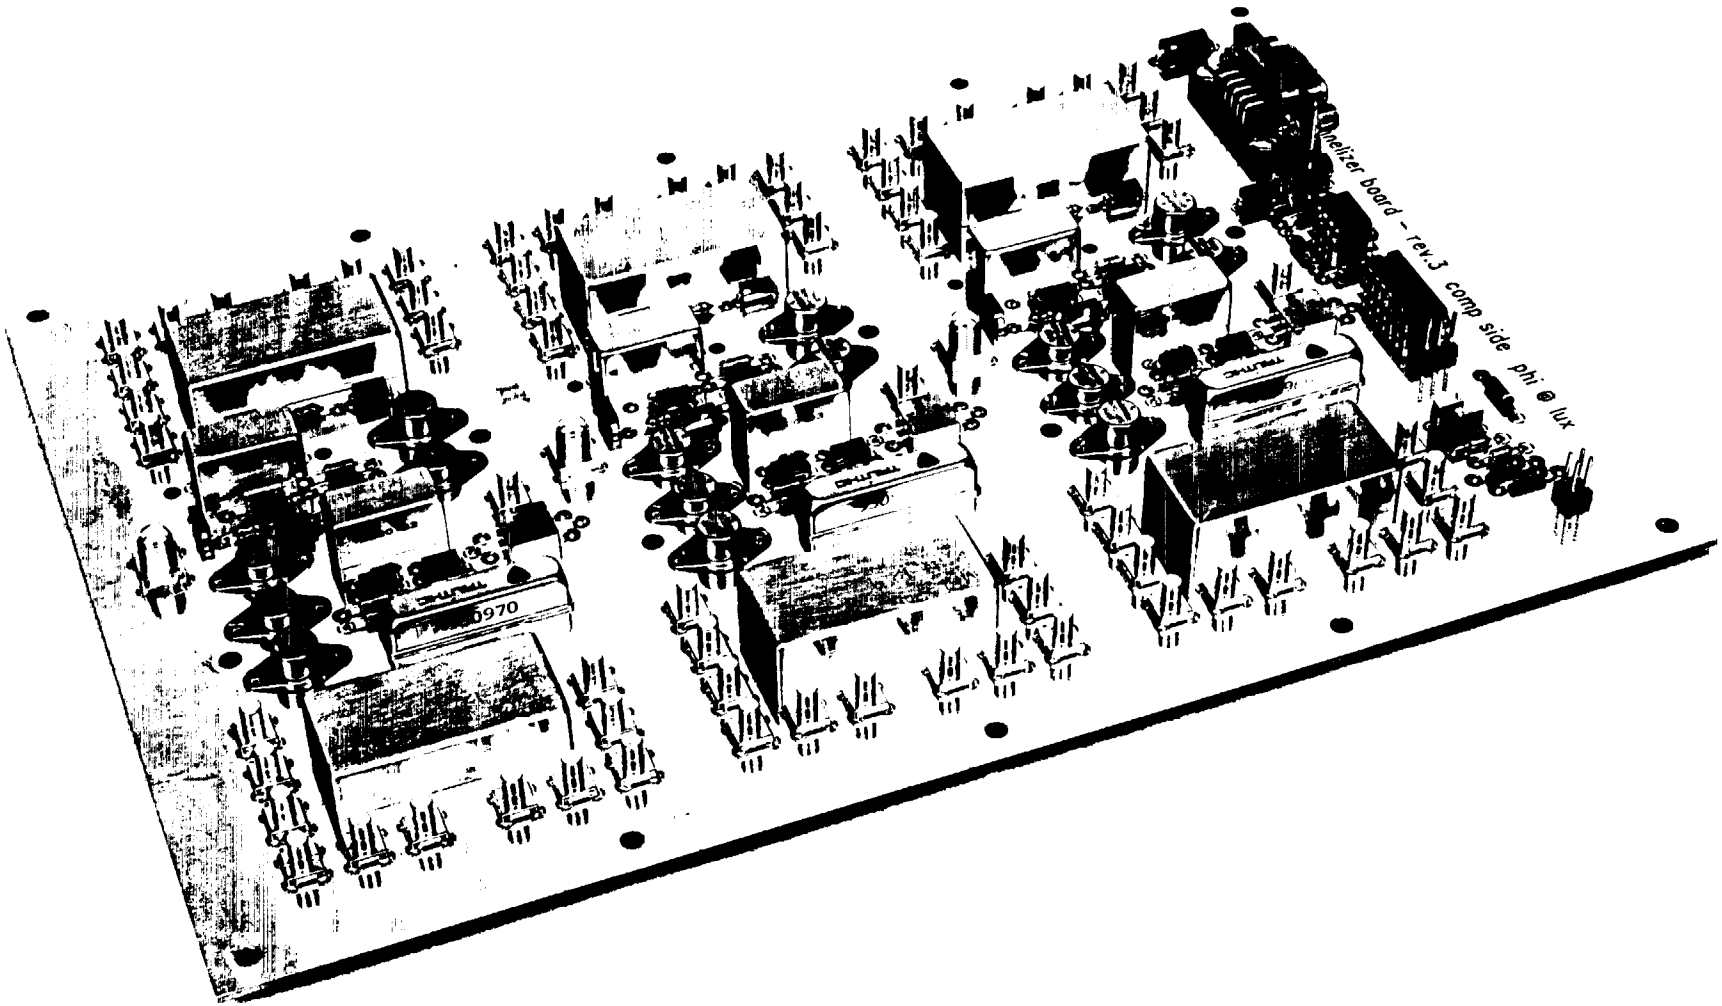

Figure 7. Completed IF Channelizer, implemented as a ground-plane printed circuit board.

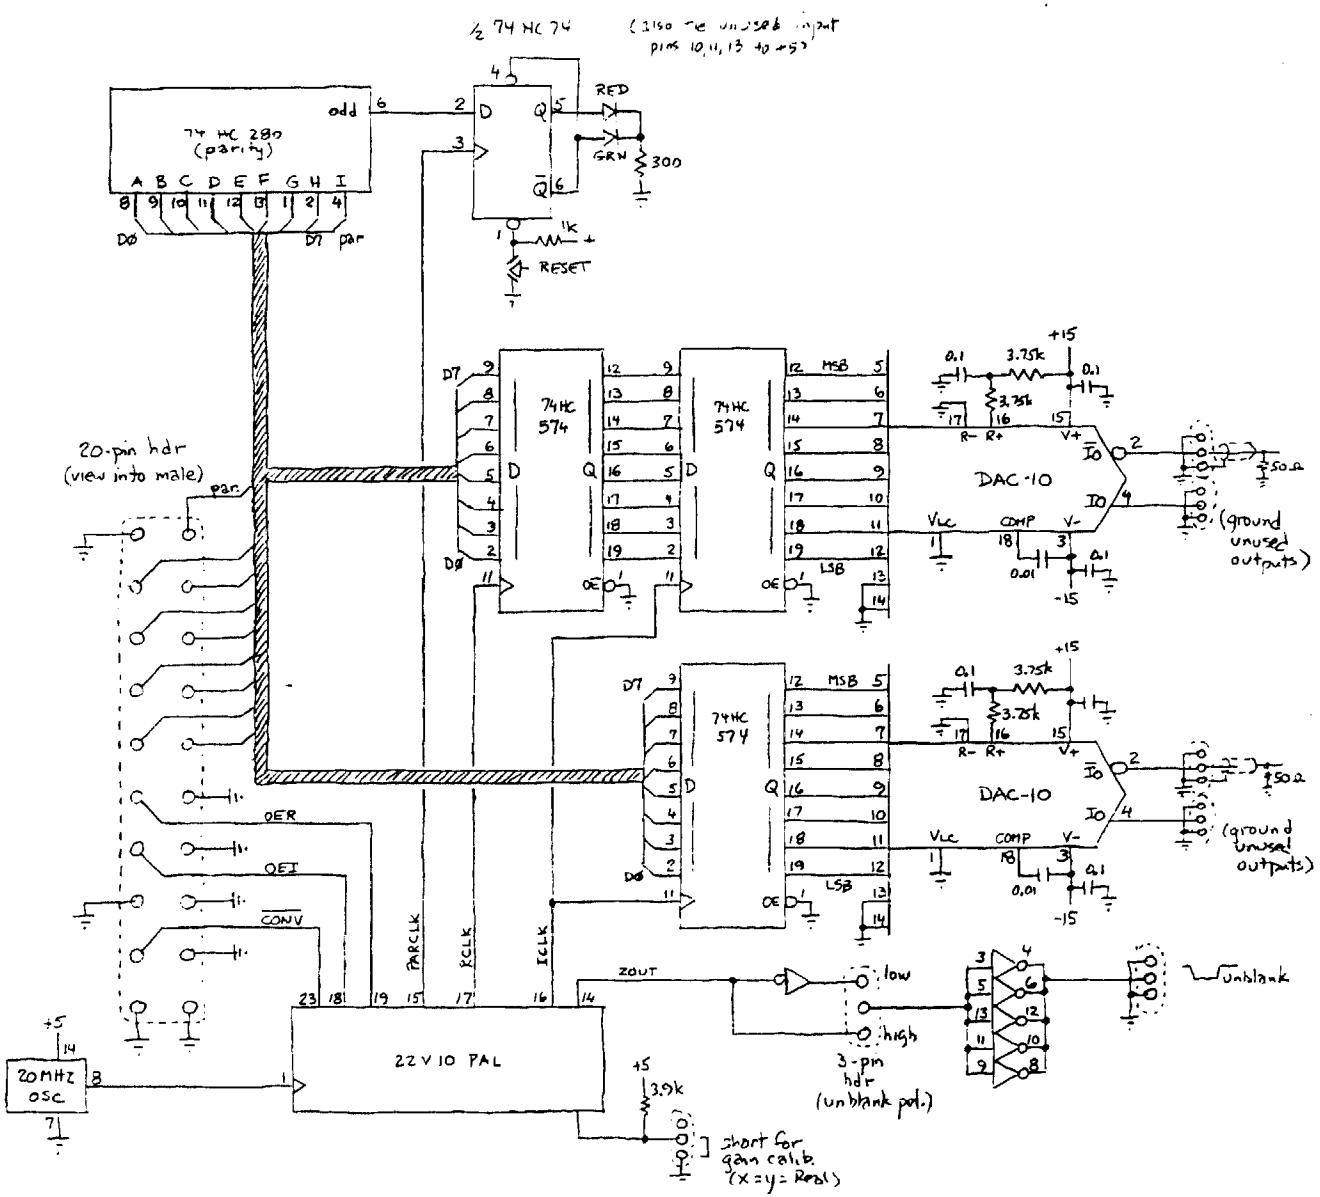

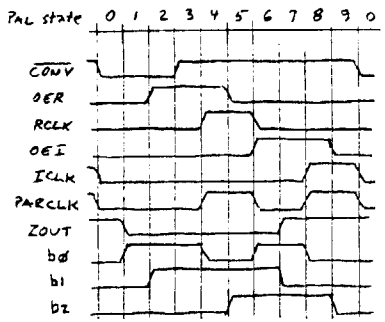

DIGITIZER TESTER		
rev	date	comments
00	12/5/93	orig. rev.

Figure 8. Tester for the mixer-filter-digitizer. The digitized quadrature baseband waveform is latched, checked for parity, and converted back to a pair of analog waveforms, along with a z-axis unblanking signal, for display on an XY oscilloscope.

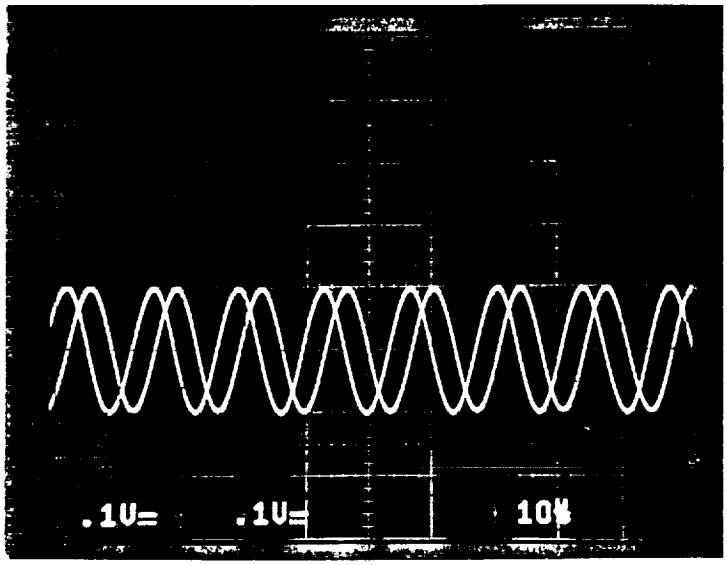

(a)

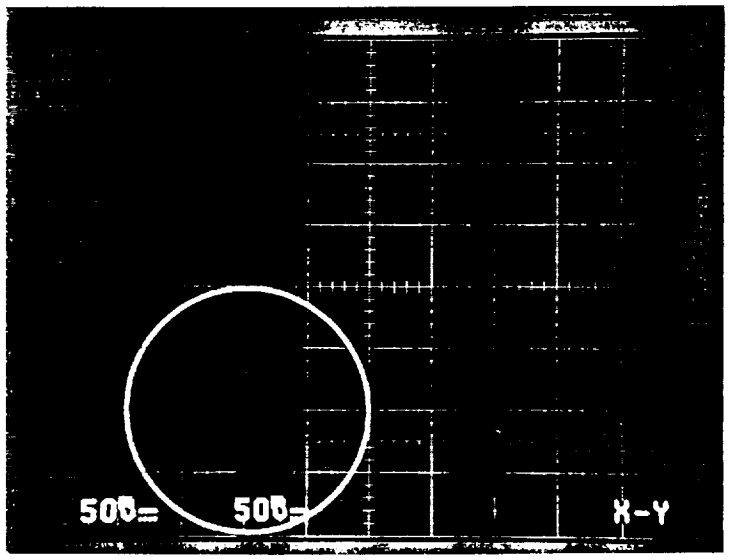

(b)

Figure 9. Analog output from the digitizer tester, for a 60 MHz LO with sinusoidal IF input offset by ≈15 kHz. The sampling rate is 2 Msps. (a) I and Q amplitudes versus time; (b) I vs Q.

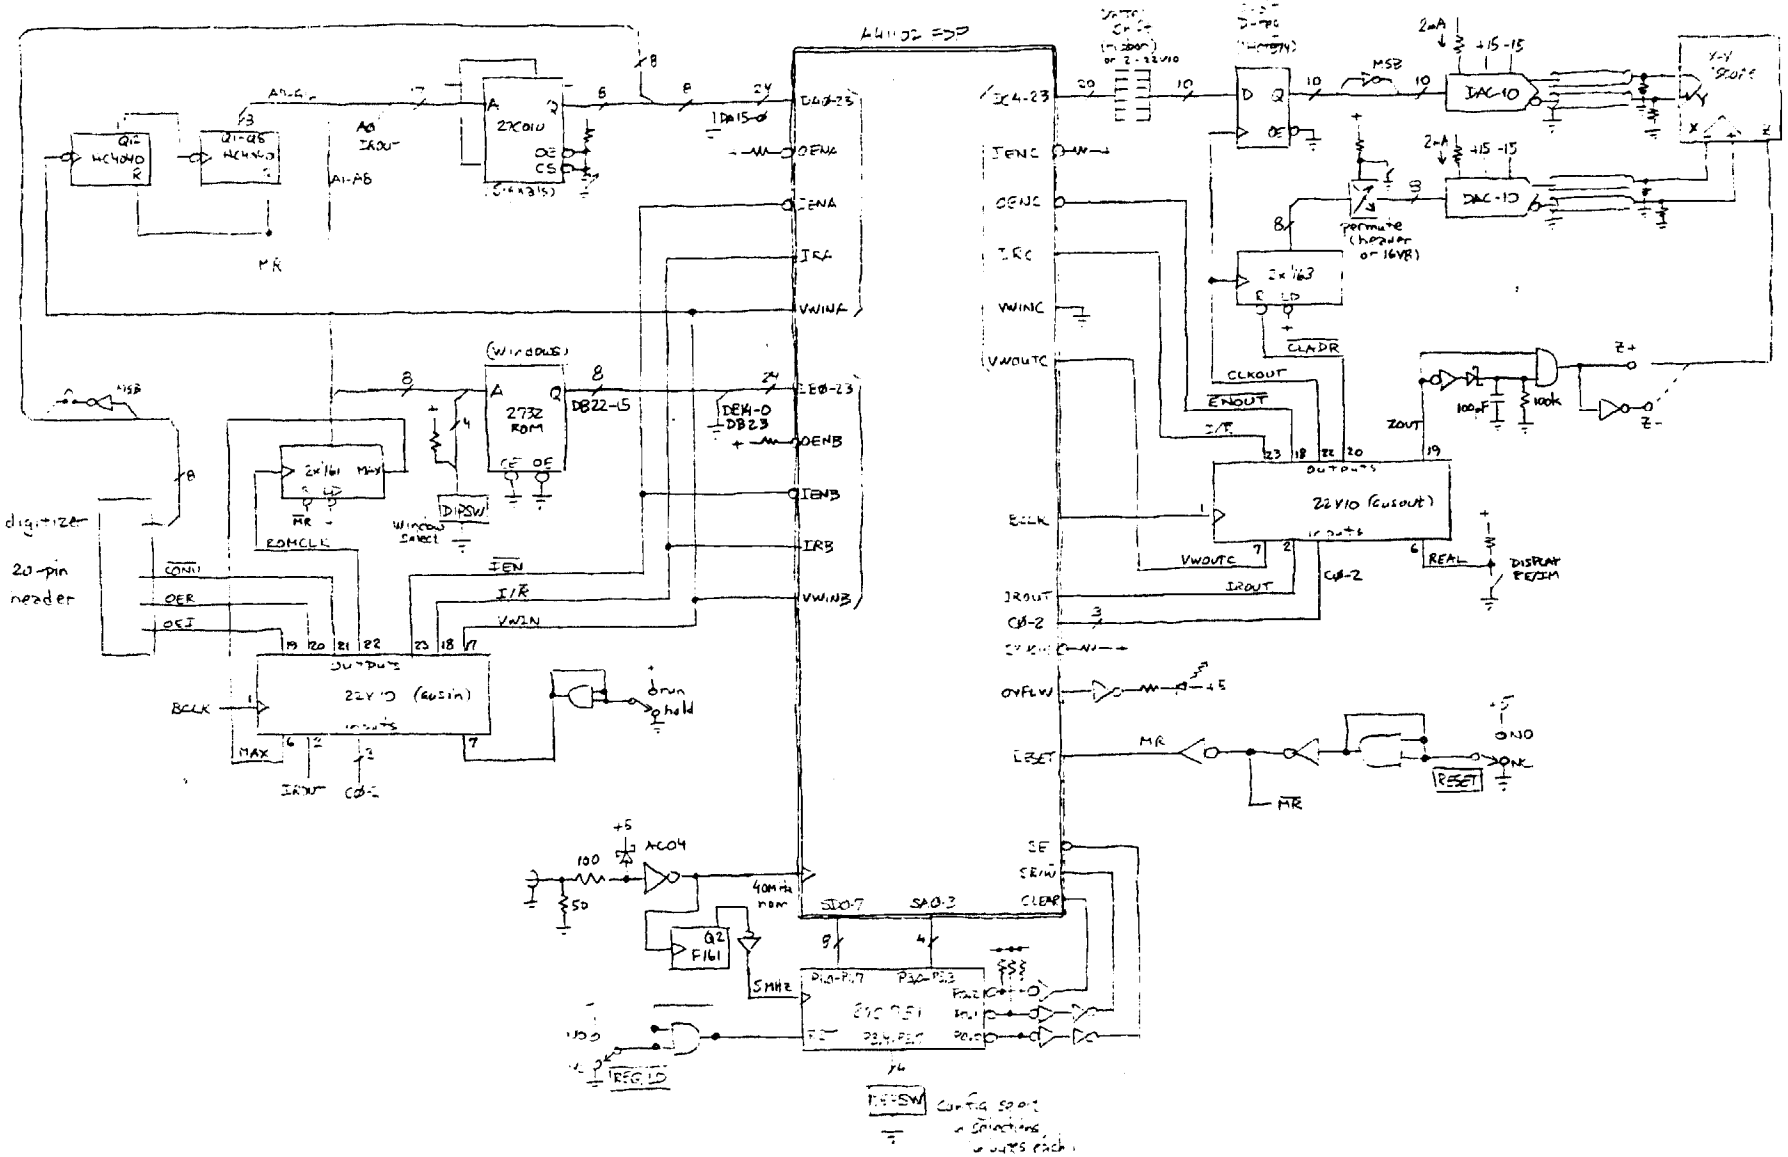

Figure 10. FFT chip tester and breadboard 256-point FFT. The input signal comes either from a ROM (containing a full set of sinusoidal basis vectors) or from the mixer-filter-digitizer. The output can be barrel-shifted and permuted, and drives an XY oscilloscope. The 87C751 microcontroller holds a set of FFT configurations in its internal ROM.

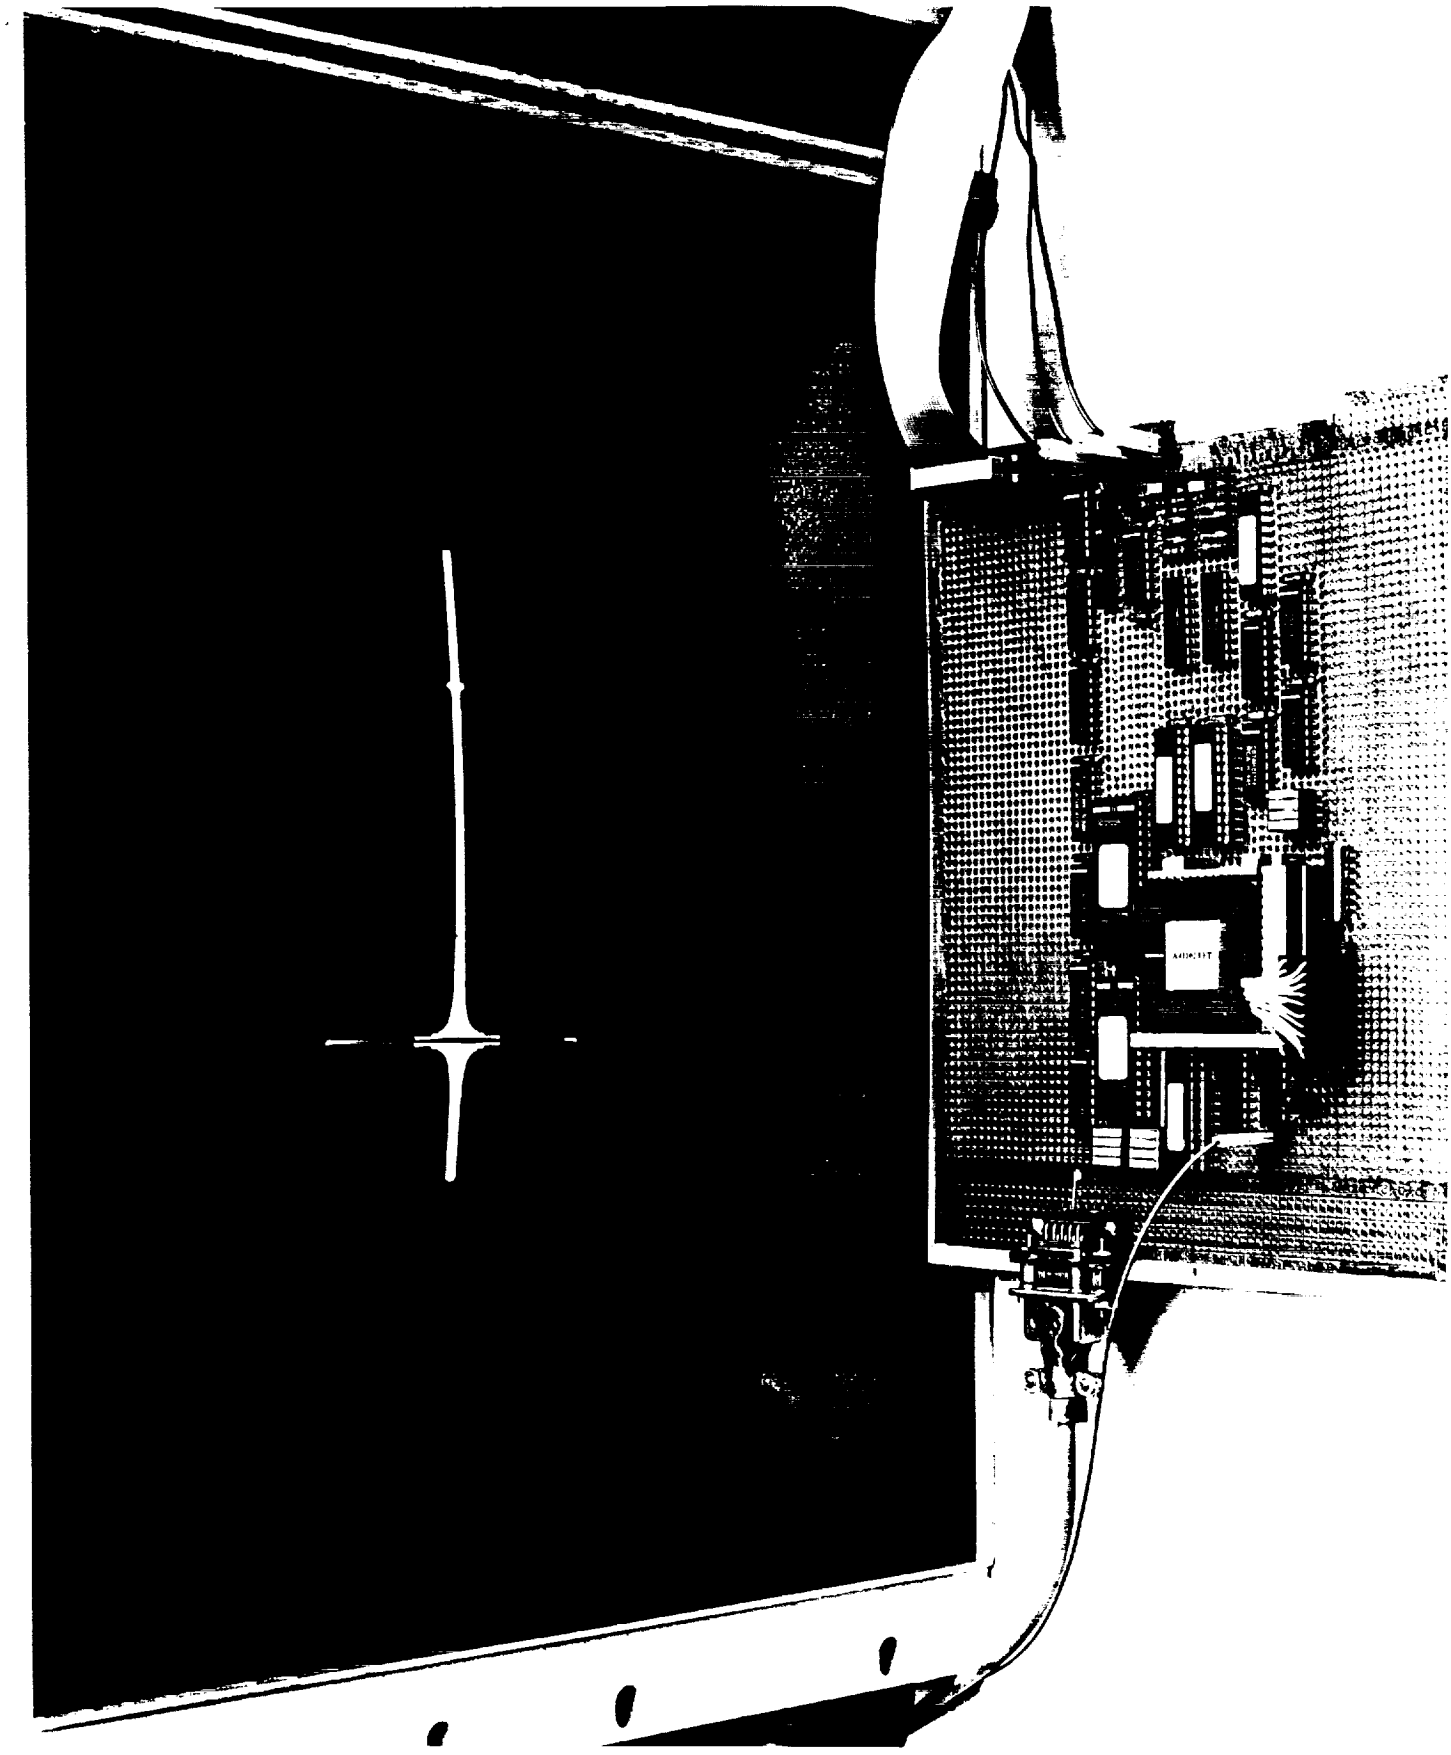

Figure 11. Photograph of wire-wrapped FFT breadboard, with real-time display of Fourier amplitudes (signed real part, linear display) plotted vs frequency. Maximum positive and negative frequencies are at the center, dc at the edges. The test signal comes from the mixer-digitizer, here chosen to lie midway between bins in order to show worst-case spectral leakage with a rectangular window. The small mirror-image response is due to imperfect baseband phase quadrature, which has not been trimmed. This prototype uses 20-bit arithmetic, and computes approximately 8000 spectra per second.

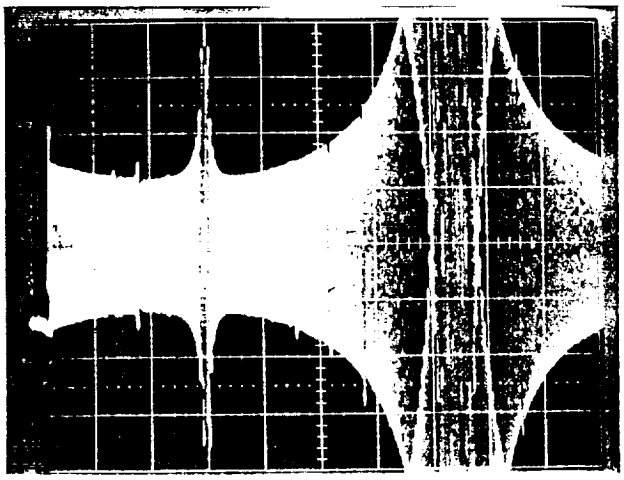

(a)

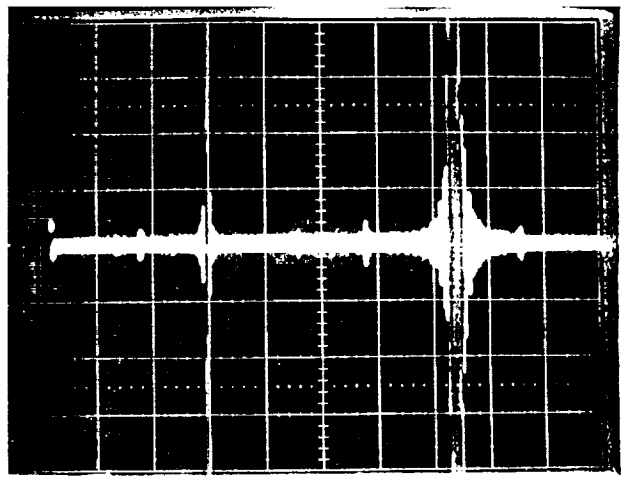

(b)

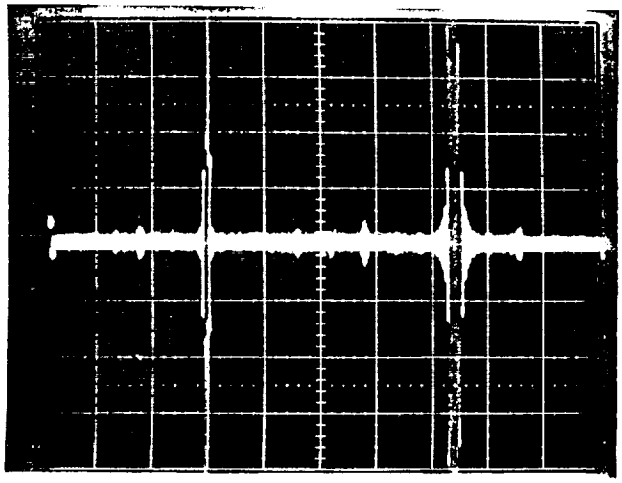

(c)

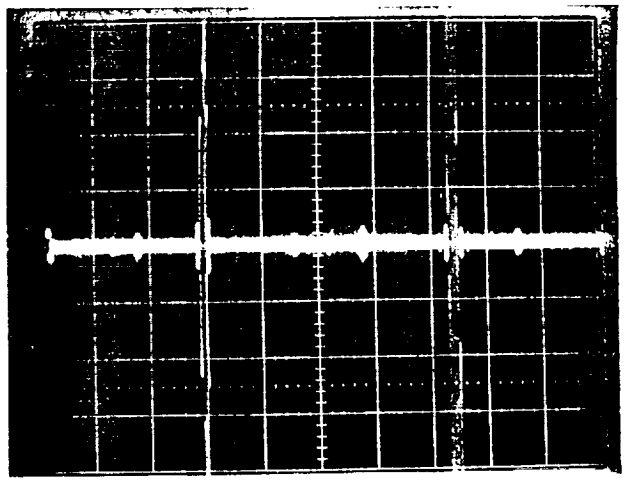

(d)

Figure 12. Effect of time windowing for suppression of low-level sidelobes ("leakage"). The test signal is a sinusoid, positioned between frequency bins for worst-case leakage. In order to make visible the low-level artifacts, the vertical scale of spectral amplitudes has been barrel shifted by 7 bits (x128, or 42dB times full-scale), producing extensive display overflow, along with full-scale display of the negative signal frequency. This *display* overflow is independent of the FFT computation, for which there is no overflow. (a) rectangular window, (b) triangular (Bartlett) window, (c) von Hann (Hanning) window, and (d) Blackman-Harris window. The latter has the highest sidelobe suppression, at the expense of main lobe broadening. The small spurious bumps are harmonic and intermodulation products, recognizable by their trajectories when the signal frequency is varied.

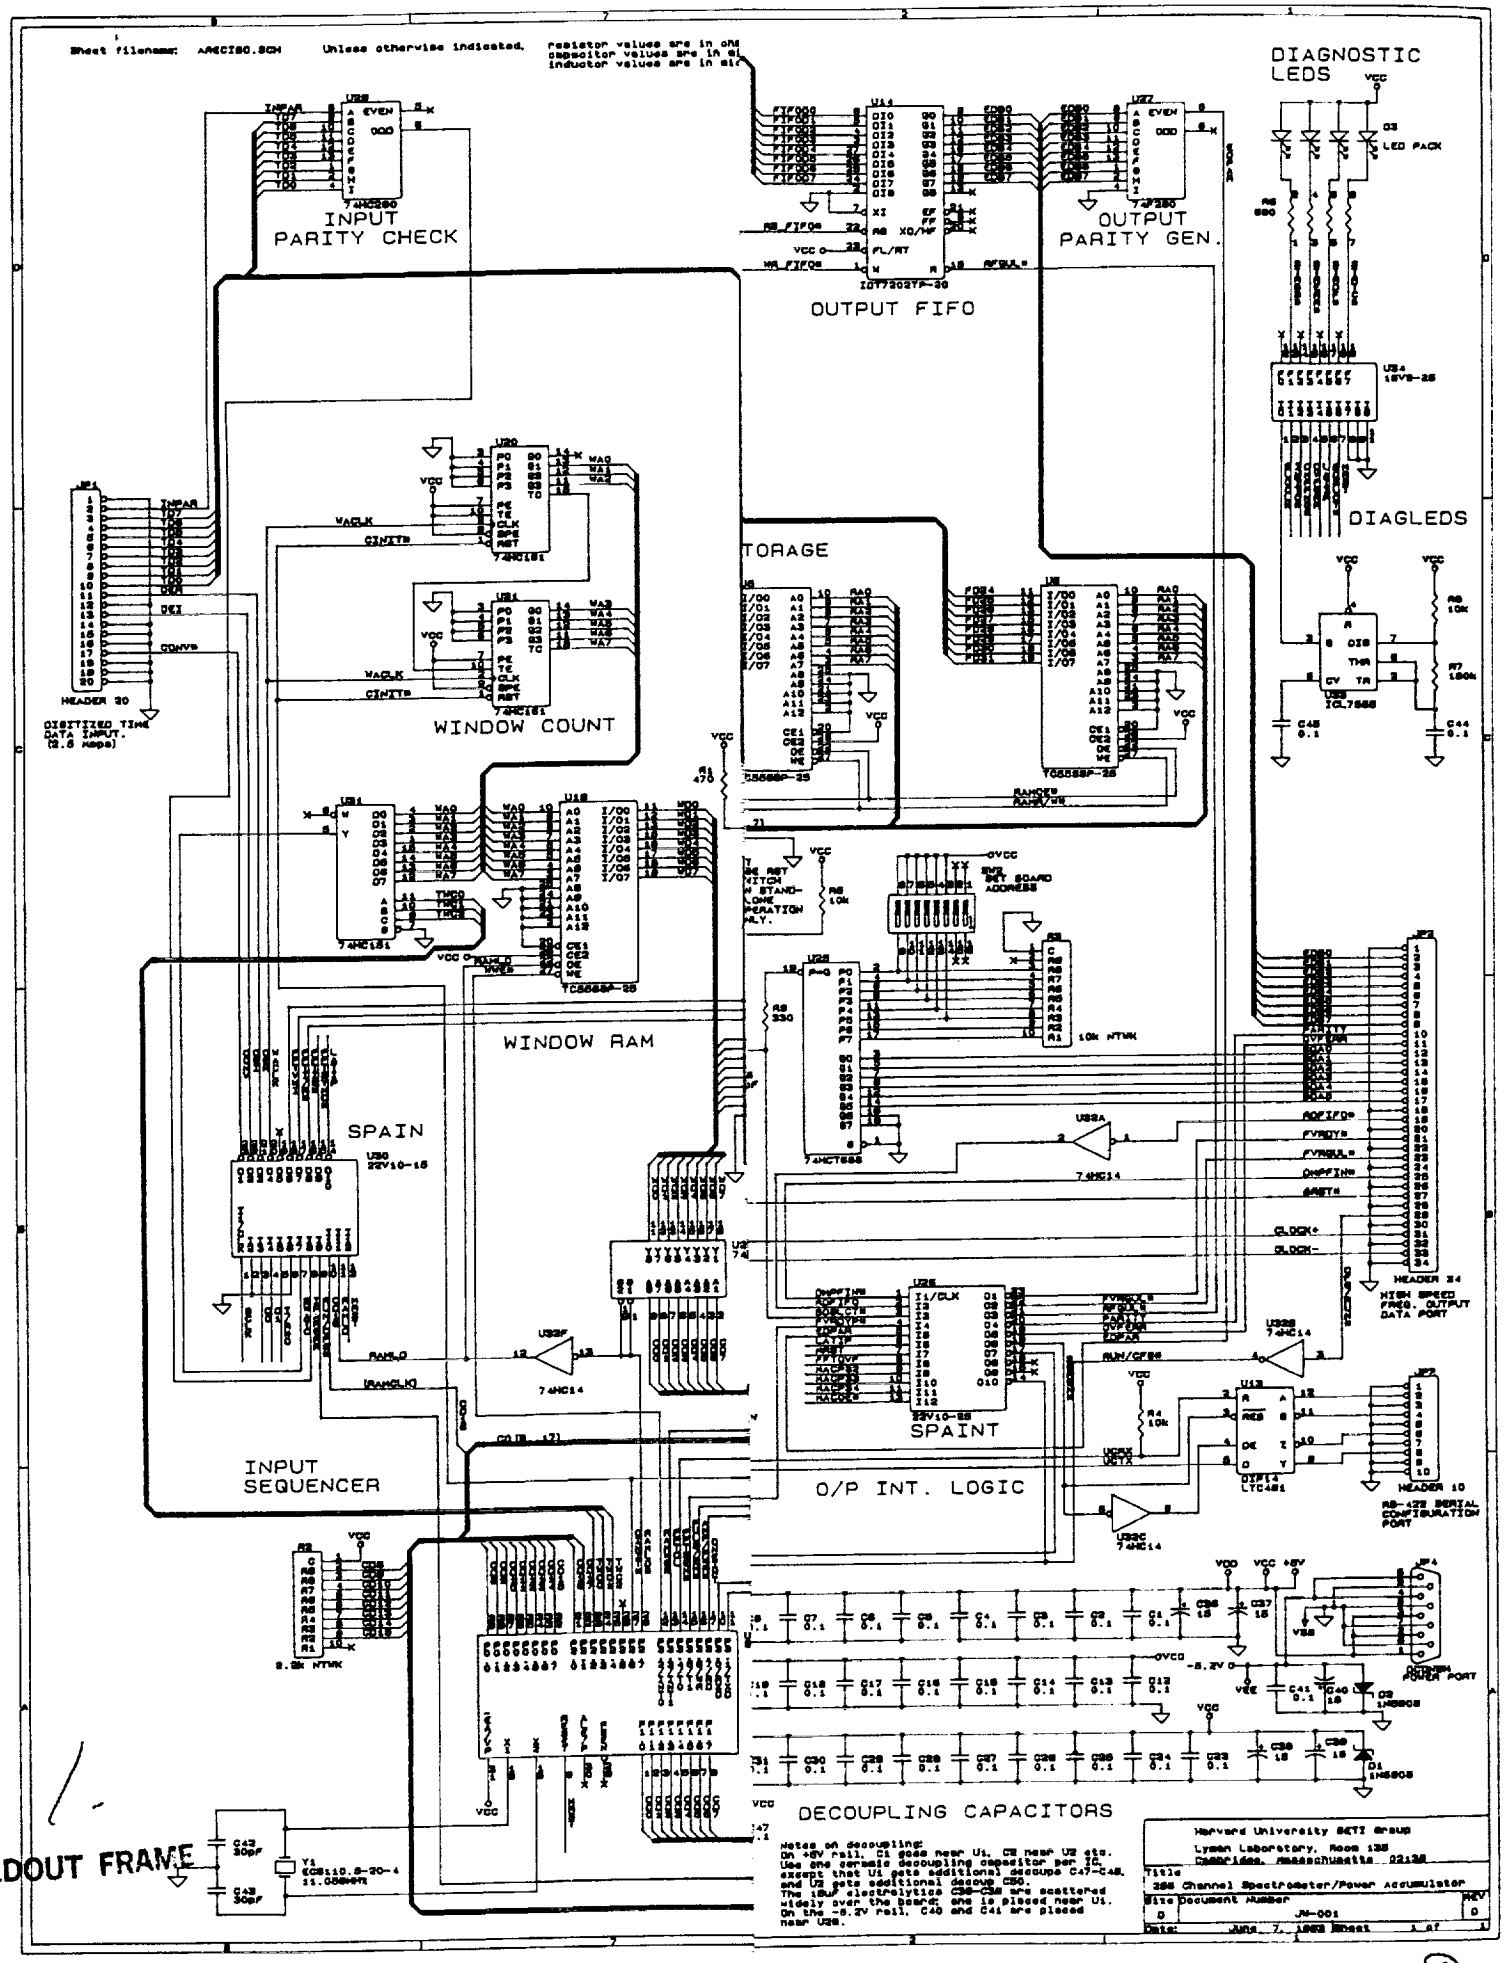

Figure 13. Spectrometer and Power Accumulator for the Arecib
can be used for long spectral integrations (up to 30 seconds, i
dynamic spectra (up to 10,000 per second, for pulsar searches and

FOLDOUT FRAME

FOLDOUT FRAME 2

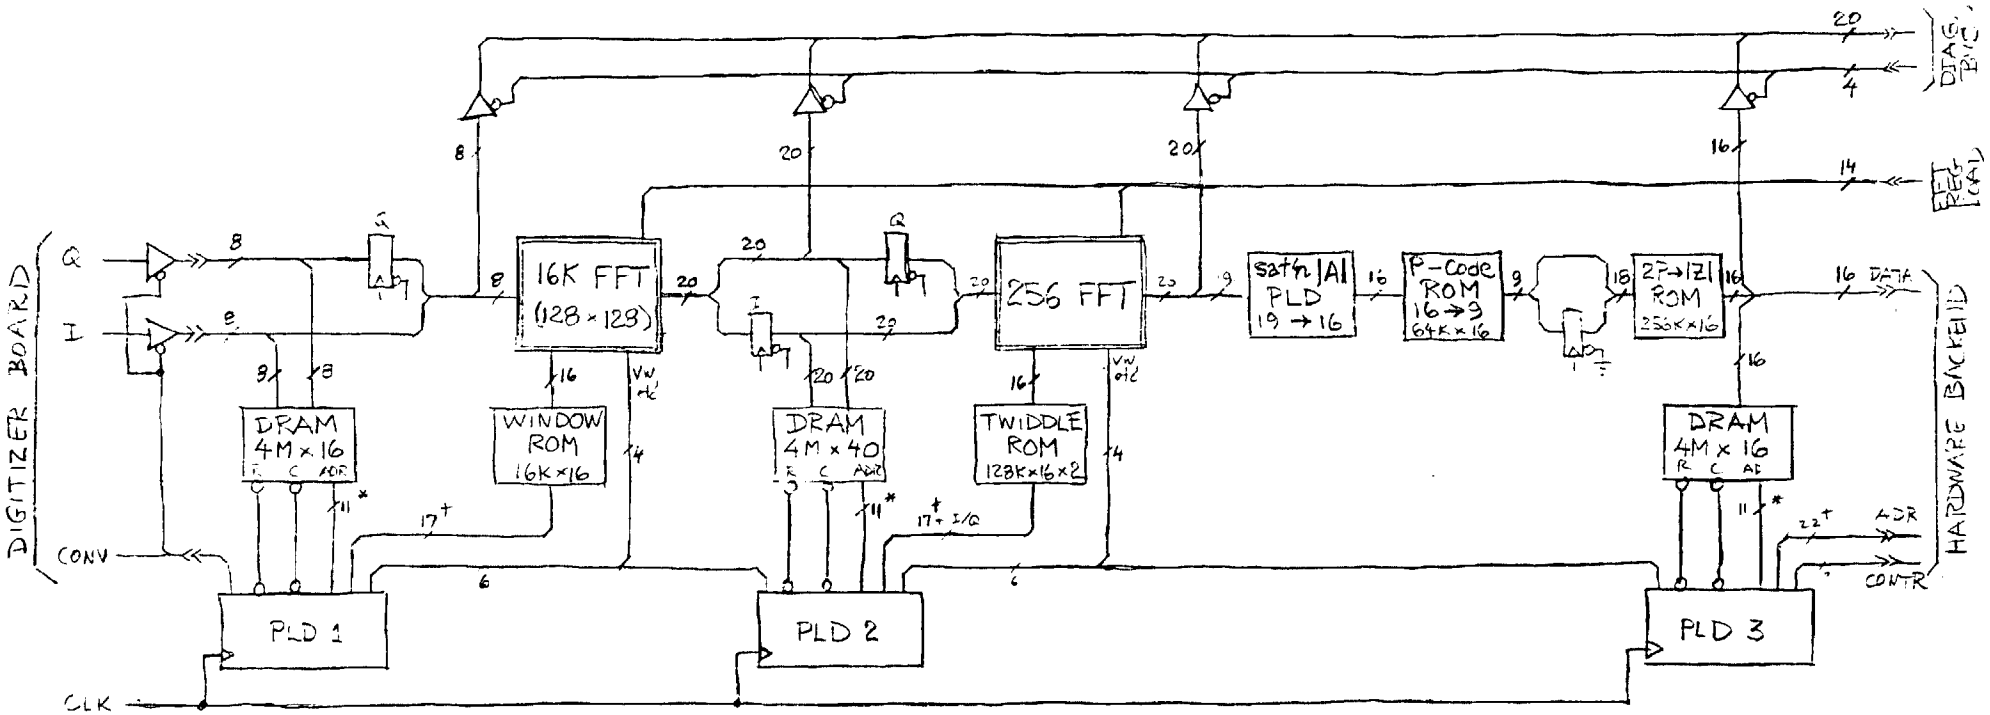

Figure 14. Block diagram of a 3-chip implementation of a 4-megachannel FFT. The final permute is needed to generate a contiguous spectrum for single-pass baseline generation and feature recognition. To save memory, a quasi-logarithmic data compression scheme is used to convert 40-bit complex amplitude data to a 16-bit magnitude format, characterized by a uniform worst-case fractional error of 0.7% (see Figure 15).

Problem: Devise an inexpensive ROM-based scheme, to compress a pair of 20-bit signed Fourier amplitudes (I and Q) into a single 16-bit unsigned Fourier magnitude ($\tilde{M} \approx (I^2+Q^2)^{1/2}$), with minimum worst-case fractional error $(M-\tilde{M})/M$.

Method: ① Begin with "saturation logic" to truncate 20-bit signed I and Q to 16-bit signed integers (a "sign de-extend"), then strip sign to generate 15-bit unsigned integers.

② Now compress the 15-bit integers (n) to 9-bit representations (ρ), minimizing worst-case fractional error upon inversion (i.e., $(\tilde{n}-n)/n$).

③ Combine the I and Q ρ-representations to generate an 18-bit address into a 256K x 16 ROM lookup table, containing the 16-bit magnitudes, namely $M_{16}(2^9\rho(I)+\rho(Q))$.

Question: What is the best way to accomplish step 2 — compressing a 15-bit unsigned integer n to a 9-bit representation ρ, with least (worst-case) fractional error?

Here are 3 possibilities ~

I. Floating Point

ρ: [e][e][e][e][f][f][f][f][f]

exp — fraction (normalized)

i.e., $n = 2^e \cdot (1fffff)_2$ Step size (worst) $= 2^{-4} = 6.3\%$

Ia. same, with "hidden bit"

ρ: [e][e][e][e][f][f][f][f][f]

exp — fraction (leading 1 omitted)

i.e., $n = 2^e \cdot (1fffff)_2$ step size (worst) $= 2^{-5} = 3.1\%$

Comment: Floating point is wasteful because only 10 of 16 exponents are used. It also has unequal step size, ranging over a factor of 2.

II. Logarithmic

$\rho = \mathrm{rnd}[\alpha \ln(n+1)]$

choose α so that $n = 0,1,\ldots,32k-1 \rightarrow \rho = 0,1,\ldots,511$ $(\Rightarrow \alpha = 49.15)$

i.e., $\rho = \mathrm{rnd}[49.15 \ln(n+1)] \leftrightarrow n = \exp(\rho/49.15) - 1$

then $n(\rho+1)/n(\rho) \rightarrow 1.021$: step size = 2.1%

Comment: This is better than floating point with hidden bit; but it is still wasteful, because of "missing codes":

n	0	1	2	4	8	16	32	64	128	256	512	1K ...	8K	16K	32K
ρ	0	34	54	79	108	139	172	205	239	273	307	341 ...	443	477	511

missing codes

III. "Hybrid" (linear, until step quantization = frac. error of logarithmic continuation)

i.e., $\rho = \begin{cases} n & n < n_0 \\ \mathrm{rnd}[\alpha \ln n - \beta] & n \geq n_0 \end{cases}$

with n_0 chosen to minimize worst-case step size (and α, β chosen to make map continuous and have proper range/domain)

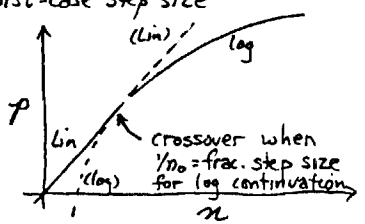

for 15-bit n
9-bit ρ,

$n_0 = 70$
$\alpha = 71.7$
$\beta = 234.6$

... and we get

n	0	1	2	...	70	71	...	100	...	128	...	1K	...	8K	16K	32K
ρ	0	1	2	...	70	71	...	96	...	113	...	262	...	411	461	511

step size (worst) = 1.4%

Figure 15. Three methods of lossy amplitude compression, and their worst-case step size upon inversion. In each case the worst-case *error* is half the worst-case step size. Conventional floating point format (with hidden bit) produces 1.6% error, a pure logarithmic format produces 1% error, and a hybrid logarithmic compression produces 0.7% error. We believe the latter to be optimal for compression of integer data, though we have been unable to find any prior use or publication. All three methods are easily implemented in a small ROM lookup table.