

SIGNAL PROCESSING IN SETI

The Search for Extraterrestrial Intelligence (SETI), now being planned at NASA, will require a prodigious amount of highly concurrent signal processing to be done in real time by special-purpose hardware.

D. K. CULLERS, IVAN R. LINSKOTT, and BERNARD M. OLIVER

Modern theories of stellar evolution suggest that almost all single stars and a substantial fraction of binary stars have planets. The Galaxy alone contains about 200 to 400 billion stars of which about five percent are "good" suns: Stars that, in their hydrogen-burning phase, radiate enough light and heat to warm an Earthlike planet and do so long enough for intelligent life to evolve. Ten billion potential life sites in the Galaxy is a reasonable estimate.

Hydrogen, carbon dioxide, methane, ammonia, hydrogen sulfide, all present in copious quantities in Earth's early volcanic atmosphere, combine under irradiation by ultraviolet light to form amino acids, polypeptides, and other building blocks of DNA. The early seas soon became a biological culture medium. It now appears that ordinary clay can polymerize these biological precursor chemicals into DNA-like structures.

Earth, fire, air, and water produced the first life on Earth and did so rather quickly, for the oldest microfossils appear in the epoch when Earth first cooled. Far from being miraculous, life seems to be forced by the chemistry of a temperate planet. In sharp contrast to the barren universe depicted 50 years ago by Sir James Jeans, we now contemplate a universe teeming with life, some of it highly evolved.

At present spacecraft speeds ($\approx 10^{-4} c$), a round-trip to the nearest star would take about 80,000 years. At 0.32 c the trip could be made in 30 years, but even with a 100 percent efficient drive, the energy required is the annihilation energy, mc^2 , of the payload. For a thousand-ton ship, this is enough to supply all U.S. energy needs for a millennium. The extreme cost of interstellar travel, even with no technological limitations save those imposed by natural law, probably explains "their" absence here and casts grave doubt on our ever reaching the stars ourselves. The probable prevalence of extraterrestrial life and its physical inaccessibility make interstellar communication seem an eventual certainty. To search for extraterrestrial intelligence by detecting signals already being radiated is a logical first step.

Electromagnetic waves are the fastest and least power consuming of all forms of radiation possible for interstellar communication. Between 1 GHz and 60 GHz the sky is quietest, so less power is required than for other frequencies. On the ground, atmospheric absorption limits this window to frequencies from 1 to 10 GHz, but since these lower frequencies are preferred for other reasons, this restriction is not considered serious (see Figure 1, p. 1153).

What sorts of signals are to be expected? Either they

"Is this, then, all that life amounts to—to stumble, almost by mistake, into a universe which was clearly not designed for life, and which, to all appearances, is either totally indifferent or definitely hostile to it, to stay clinging on to a fragment of a grain of sand until we are frozen off, to strut our tiny hour on our tiny stage with the knowledge that our aspirations are all doomed to final frustration, and that our achievements must perish with our race, leaving the universe as though we had never been?"

Sir James Jeans

The Mysterious Universe (1933)

are not cooperating and we must detect signals radiated for their own use (the "eavesdropping" mode), or they are deliberately trying to attract our attention (the "beacon" mode). In the first case, we might hope to find carriers such as are radiated by our UHF television stations, or as might be radiated by orbiting solar power stations. We have no hope, without beaming on their part or much larger receiving antennas on our part, or both, of receiving complex modulation; rather we will be fortunate to detect signals whose spectra are no wider than required by their duration.

Beacons are more likely to transmit pulses than continuous-wave (CW) carrier signals, since less average power is required. Although the power may be higher and the received signal-to-noise ratio (SNR) therefore greater while the pulse is on, no power is needed for the long off periods. The principle is the same as the sweeping lighthouse beam or the flashing strobe light: Both produce attention-getting flashes at a low *average* power.

Because the detectability of a pulse in the presence of noise depends only on the total energy, microsecond-long pulses would need a million times more peak power than second-long pulses. We are therefore inclined to expect rather long pulses. Unfortunately there does not appear to be any "natural" length, so we are faced with the task of providing nearly matched filters for pulses anywhere from a few seconds to a few milliseconds in length.

Both CW and pulse detection involve sifting through hundreds, even thousands, of megahertz of spectrum for anomalies that may be only a fraction of a hertz wide. Only recently, with VLSI technology, has it become possible to build at a reasonable cost the digital data-processing equipment needed.

The search strategy of the NASA program uses two distinct modes: an all-sky survey and a targeted search. The all-sky survey will cover more spectrum at a lower resolution (32 Hz) than the targeted search. Because of this and because the effective pointing time in any direction is so short, the sky survey is much less sensitive ($\sim 10^{-24}$ W/m²). It may, however, discover powerful CW signals coming from unexpected directions.

The targeted search assumes that good stars are rather like ours. It will examine some 800 solar-type stars out to about 80 light-years, spending about 20

minutes on each star for each frequency band. It will cover the spectrum from 1 to 3 GHz (with emphasis on the water hole) to a resolution as narrow as 1 Hz. The sensitivity will be as high as 10^{-27} W/m² for CW signals and $\sim 10^{-28}$ W/m² (average) for pulses.

All presently planned searches will use existing radio-astronomy antennas. If no ETI signals are discovered, our society will have to decide whether SETI justifies a dedicated facility of much greater collecting area.

MULTI-CHANNEL SPECTRUM ANALYZER

Like the spectroscopist looking for sharp spectral lines in a broad spectrum, we are searching for very narrow band anomalies in a broad radio-frequency band, so the obvious first tool we need is a radio-frequency spectrometer. This cannot be the usual swept receiver type of spectrum analyzer, for these look at any signal only a small fraction of the time, lose most of the signal energy, and hence degrade the SNR. It could be a multi-million point digital Fourier transform (DFT) computer operating in real time. Our actual choice is the multi-channel spectrum analyzer (MCSA) shown in symbolic form in Figure 2 (p. 1154).

A 10-MHz band of the IF signal is converted to base-band by two mixers operating in quadrature. The two resultant 5-MHz signals are sampled at 10.6 MHz to form two four-bit data streams representing the real and imaginary parts of the input signal. Four-bit (16-level) quantization was deemed adequate since the loss in SNR is no more than 0.05 dB. The digitized complex data stream is then processed by two digital filter banks in tandem followed by multiplexed DFTs executed by microprocessors.

The first band-pass filter is implemented as two identical filter banks, one supplied directly with the input signal and the other with the input signal shifted in frequency by 73.7 kHz. The shift is done by multiplying the input data by a complex sinusoid. Thus one bank produces what will become the even-numbered channels, and the other produces the odd. Each bank produces 72 channels each 147.4 kHz wide, of which only the central 73.7 kHz is used by band-pass filter number two, and is free of aliasing. Of the 144 channels, 16 at each end lie in the transition region of the analog filter and therefore contain noise and possibly signals aliased

from outside the band of the MCSA. In effect the analog signal is oversampled by 9 : 7, and the digital signal is oversampled 2 : 1 in band-pass filter one.


The central 73.7 kHz of each of the 112 selected bands is split by band-pass filter number two into 72 bands, for total of 8064 bands, each 1152 Hz wide; with their transition regions overlapped. The central 1024 Hz of these bands is free of aliasing and is contiguous band to band. In effect there has been an oversampling of 9 : 8 in band-pass filter number two.

The 1152-Hz bands are then subdivided by the multiplexed microprocessors DFTs into 36 bands 32 Hz wide

(4 of which are discarded) or 1152 bands 1 Hz wide (128 of which are discarded).

Both band-pass filters consist of finite impulse response (FIR) filters followed by 72-point inverse DFTs. The "zero-frequency" output channel of the DFT is simply the convolution of the input signal with the FIR. The passband, centered on d.c., is the transform of the FIR. All the other channels have passbands of the same shape, and the effect of the DFT is to shift them all to baseband.

The FIR and equivalent low pass filter characteristic of band-pass filter number two are shown in Figures 3a



Synchrotron radiation from electrons spiraling in the magnetic field of the Galaxy produces a noise that rises rapidly with decreasing frequency below 1 GHz. For frequencies above 60 GHz, the quantum nature of radiation produces a noise proportional to frequency. In between, the only noise is the cosmic background radiation, a relic of the big bang. In this silent valley, we may discover signals radiated by extrater-

restrial intelligent life. The low-frequency end of the window, best for many reasons, is marked by the hydrogen (H) and hydroxyl (OH) spectral lines, symbolic of water and therefore of life. The 300 MHz of spectrum between these signposts is called the water hole. Here we may meet other species as happens at more mundane water holes.

FIGURE 1. The Microwave Window




FIGURE 2. Multi-Channel Spectrum Analyzer

Two finite impulse-response digital filters in tandem resolve over 8 MHz of spectrum into 8,064 channels 1,024 Hz wide. These are then subdivided into channels as narrow as 1 Hz by multiplexed microprocessors executing digital Fourier transforms. The algorithms are stored in random-access memory down-loaded from the system computer disk file. Thus the number of channels and the resolution can be changed as desired without altering the hardware. Present proof-of-principle prototype employs at least one of each type of circuit involved and has 74,000 channels. For reasons of size and economy, the full bandwidth analyzer will incorporate very large scale integration.

and b. The FIR is composed of 1728 coefficients, 144 between each pair of numbered taps. (Because of the symmetry, only 864 values need to be stored.) It is apparent from Figure 3b that, in order to get flat response to the edge of the passband and full rejection of all


aliased frequencies, the transition region should be half as wide. This would require 24 taps and 1728 stored coefficients. Whether to do this or to oversample by $\frac{5}{4}$ rather than $\frac{3}{4}$ remains to be decided.

The MCSA architecture described above was chosen

over a more conventional monolithic eight-million-point DFT processor for several reasons. First, a proof-of-concept prototype could be (and has been) constructed without providing all eight million channels, merely by building one board of each type. This substantially reduced the hardware cost of the prototype. Second, the full machine would be highly modular with a great deal of parallel processing. This meant that the number of different boards was reduced, simplifying the stocking of spares. Further, a failure in any downstream board would wipe out only a small portion of the spectrum, and could easily be located and replaced. Third, the high degree of isolation between the

bands of the band-pass filter was felt to be important when operating in the presence of radio-frequency interference. Finally, the operating binwidths can be changed simply by changing the microcode that controls band-pass filter number two and the final DFTs.

All outputs of the MCSA are available either as complex amplitudes or as the sum of the squares, which is the detected power. The MCSA also allows the average power per channel to be computed and a threshold to be set at some arbitrary multiple of this average power. The computer can then ask only for the power and address of those bins that exceed threshold, thus greatly reducing the amount of data to be transferred.



The FIR for the present band-pass filter number 2 is shown in (a), while the resulting equivalent low pass frequency response is shown in (b). The passband extends from $-f_0$ to f_0 . The sampling frequency is at $f_s = 18f_0/8$ so the spectrum is folded at $9f_0/8$ (and $-9f_0/8$). The frequencies above $10f_0/8$ (and below $-10f_0/8$) alias into the passband. The aliased response is 21 dB down at f_0 and drops to over 60 dB down at $7f_0/18$. To prevent any aliasing above the 60-dB level will require about twice as long an impulse response.

FIGURE 3. Band-Pass Filter Characteristics

CW SIGNAL DETECTION

At first glance the detection of a CW carrier seems very simple. Since the matched filter for a monochromatic signal has zero bandwidth, subdivide the spectrum into the narrowest practical channels and, after integrating for the observing time, look for excess power in one of the channels. In other words, equip the MCSA with an accumulator and report out any bins in which the accumulated energy exceeds a certain threshold, chosen so that the false-alarm rate is tolerable.

Unfortunately very few signals are sufficiently monochromatic for this simple scheme to work. Even if the transmitter generates a monochromatic signal, acceleration along the line of sight of the transmitter or receiver, or both, due to planetary rotation and orbital revolution can introduce drift rates as high as ~ 1.0 Hz/s at carrier frequencies of 1.5 GHz. We can of course correct for the Doppler drift due to our own acceleration, and if the signal is being beamed at us, they can do likewise. Thus in the case of a nearby CW beacon we might expect a drift-free signal. But in general we should be prepared to detect drifting signals.

When the signal straddles two bins, the power in each of the two bins is only 40 percent of that in a single bin when centered. A slowly drifting signal will in general be near bin center in some samples and will straddle two bins in others. This produces a loss in detectability that is clearly evident in Figure 4. We are currently studying ways of avoiding or reducing this loss such as by providing overlapping bins, or "interbins."

Neglecting this problem, it is clear that *patterns* of straight lines are visible in Figure 3 when no signal is evident in a single spectrum sample. How shall we automate the detection of these straight-line patterns? The Cyclops study (1971) proposed recording all spectra during the observing time and then adding these with amounts of "skew" corresponding to all distinguishable drift rates up to ± 1 bin per sample. Two drift rates are distinguishable if their total drift in the observation time differs by 1 bin (or more). Thus in n samples, there are $2n - 1$ distinguishable drift rates. For the skew that matched the actual drift rate, all the energy would accumulate in a single bin.

This is an excellent method, but it is memory and computation intensive. Each bin requires about a byte of memory so a full spectrum register has $N \approx 10^7$ bytes. With $n = 10^3$ samples per observation, $nN = 10^{10}$ bytes are needed (twice as much if interbins are used). These 10^{10} numbers must be added in $2n - 1 \approx 2 \times 10^3$ ways. For 1-Hz bins, this means 2×10^{10} additions per s.

We are currently exploring a pruning algorithm that saves memory and computation by carrying out the accumulation in three phases. In the first phase, a few spectra, say 20, are stored in as many full spectrum registers and are tested for excess energy along all distinguishable lines: $N(2n - 1)$. All those, say $1/10$, that pass a lenient criterion are retained.

In the next phase, which is longer, each of the qualifying lines of phase one is allotted some memory, for a total of about $1/10$ of a full spectrum register. Samples are now dedrifted before storing, using the drift rates found in phase one, in order to confine the signals to the assigned blocks of memory. To store, say, 200 samples again requires 20 full spectrum registers. All lines are again tested against a less lenient criterion. The possible spread in drift rates is now only $1/10$ as great.

In the third phase, only one percent of a full spectrum register is needed per sample for the surviving candidates, so in principle we have enough memory for 2000 samples. In the calculations made to date, 800 were used. The drift rates are further subdivided, and the number allowed to pass is that needed to keep the overall false-alarm rates acceptable.


Preliminary calculations indicate that the detection sensitivity of this as yet unoptimized process is within about 1 dB of the Cyclops method, but the memory and computation required are about $1/50$ as great. The strategy is to use a very low initial threshold and then weed out the false alarms that this produces.

If we can tolerate a false alarm in a fraction k of the observations and if there are N bins per spectrum and the spectra are added $2n - 1$ ways, then the allowable false-alarm probability per addition is $k/N(2n - 1)$, or about 10^{-12} for $k = .02$, $N = 10^7$, and $n = 10^3$. As the curve in Figure 5 (pp. 1158–1159) for a false-alarm probability of 10^{-12} shows, the SNR needed to have a 50-50 chance of detecting the signal is 14.4 dB for a single sample ($n = 1$), but only -0.6 dB when 100 samples are averaged, an improvement of 15 dB. But now suppose that all the energy received in those 100 samples had been sent as a pulse in only 1 sample time. The SNR would then be -0.6 dB + 20 dB or 19.4 dB rather than the 14.4 dB needed: a clear gain of 5 dB. This is part of the gain in energy efficiency obtainable with pulses over CW signals and is what leads us to expect pulsatile signals from beacons.

PULSE DETECTION

The fundamental assumptions we make in pulse detection are that the pulse-to-pulse time interval is constant and that any drift in frequency is comparable to that for a CW signal in the same time. In other words, we confine our attention to regular pulse trains (with perhaps one or two pulses missing) lying on nearly straight lines in a frequency versus time plot. In a regular pulse train, the individual pulses can have a fairly high probability ($\sim 10^{-2}$) of having been caused by noise alone; all we require is that the *joint* probability of the entire train being due to noise alone be low enough ($\sim 10^{-12}$) to provide a tolerable false-alarm rate.

The pulse-detection algorithm has the following basic steps: Set a threshold such that the number of hits per spectrum sample (i.e., bins that exceed the threshold) is as great as can be handled by the (special-purpose) processor. Record the bin address and power of each hit. Do this for each spectrum sample. Form all



Each line of raster shows the power in each output channel, or "bin," for 120 adjacent channels during one sample.

Ninety successive samples are shown. There is a general field of noise in which the probability of a given brightness falls off exponentially with brightness. Included are three drifting CW signals. The two with a slope of 45° represent signals decreasing in frequency 1 bin per sample. The right-

hand signal is always centered in the bin at the middle of the sample. The left-hand signal straddles two bins at the middle of the sample. The resulting loss in detectability is evident. The third signal is increasing in frequency $\frac{1}{2}$ bin per sample and is alternately at bin center and at bin edge. For all the signals, the SNR is 2; that is, the signal power is twice the average noise power per bin.

FIGURE 4. Detection of CW Signals

lines consisting of a hit in one spectrum sample and a hit in a later sample that do not exceed the maximum drift rate. Predict the sample and bin numbers for the third and fourth pulses in any of these putative pulse trains. If a third or fourth pulse occurs, retain that line; if not, discard it. After the observation period, compute the joint probability of all retained lines being due to

noise. If any of these is low enough, announce the preliminary detection and repeat the observation.

After 100 spectra the field of hits in 500 adjacent bins might look like Figure 6 (p. 1159). Finding a regular pulse train in this field by eye takes a long time, but a VAX/750 using our present algorithms in software can find any pattern in 1,000 spectra of 74,000 bins each in




FIGURE 5. Signal-Detection Statistics (left)

Blue curves show the threshold in dB above mean noise power needed to achieve the indicated false-alarm probability per bin as a function of the number, n , of samples accumulated, or averaged. Red curves show, as a function of n , the SNR needed to cause the sum of signal and noise to exceed the threshold (thus detecting a signal) with 50 percent probability. For high SNRs the decrease in required SNR is nearly 3 dB per doubling of n . For low SNRs the improvement is about half as great. Basically this is because of the poor performance of a square law detector at low SNRs. Pulses having the same average power as a CW signal are more detectable because the SNR is higher when they are present.

real time. To handle 10^7 bins will require special-purpose hardware. (Try to find the regular pulse train before looking at Figure 8.)

Two practical problems complicate the above simple approach. The first is that we have no a priori knowledge of the pulse length. We cannot therefore choose a single binwidth in the MCSA that provides a matched filter for the pulse. Instead it appears that we may have to provide binwidths in octave steps from, say, 1 Hz to 32 Hz or even 1024 Hz. The second is that we cannot expect each successive pulse to lie centered in frequency in a bin or to coincide in time with a spectrum sample.

The relative sensitivity of two algorithms can be determined by comparing them under the same condi-




The threshold has been set so that noise alone will exceed it 1 percent of the time, producing a "hit." The hits in 120 adjacent bins are shown for 90 successive samples. In this random field of hits, a regularly recurring set of hits produced

by a train of pulses 8 dB above the mean noise power and drifting less than 1 bin per sample has been added. Look carefully. Can you find it? Compare your answer with the computer's, shown in Figure 7.

FIGURE 6. A Field of Hits

tions using computer simulation and a Monte Carlo process. The algorithms must give the same false-alarm rate with noise alone. The algorithm that detects the signal more often for the signal parameters chosen is the more sensitive. Figure 7 shows the results of such a comparison for three types of signals. The pulse detector with synchronous sampling shows significantly better performance than the other two detectors. Its performance in the regime plotted is, in fact, almost as good as a synchronous detector. This is because the pulse energy is confined to a single bin per pulse and the (peak) SNR is rather high. At high SNRs all detectors approach the performance of a synchronous detector.

However, the loss due to lack of pulse-sample syn-



Comparison, by Monte Carlo simulation, of the performance of CW detection and two cases of pulse detection: one in which the pulses comprising the four-pulse train are all synchronous with the time samples that generate the spectra, and the other in which the pulses all straddle the successive time samples. The scattering of the energy in the latter case degrades the performance by almost 6 dB. The curves assume a 1-Hz bin width and 1000 spectral samples. Notice how rapidly the probability of missing the signal drops with increasing SNR.

FIGURE 7. Detector Performances

chronization is serious; one hates to throw away $\frac{3}{4}$ of the antenna area. It appears that just as interbins are desirable for CW detection so both interbins and intersamples are desirable for pulse detection. To provide these as well as octave binwidths will approximately double the computing needed. Various alternative approaches are currently under study.

THE PATTERN DETECTOR

The detection of either drifting CW signals or pulse trains is computation intensive and will require special-purpose hardware if it is to be done in real time over 10 MHz of bandwidth or more. The design of this hardware is presently in a very preliminary stage so only general remarks can be given here. It is apparent that, since the same tests must be applied over a wide frequency range with each test involving a relatively narrow span of frequencies, the problem naturally lends itself to highly concurrent processing: 10 MHz of spectrum might require 100 or more microprocessors each serving 100 kHz of spectrum or less.

If the spectral bands served by the processors are contiguous, signals near the band edge of one processor may drift into the band of the adjacent processor during the observation time. To avoid losing track of such signals, one must either anticipate the event and transfer control to the adjacent processor or make memory addresses outside the nominal band available to each processor. The simplest solution seems to be to overlap the bands of the processors (by about 10 percent) so that all signals within the maximum drift rate will remain within the band of at least one processor.

Two polarizations, left-hand circular (LHC) and right-hand circular (RHC), should be received simultaneously by the SETI receiver. Circular polarization remains circular after propagating through the interstellar medium. Thus, if a transmitter radiates either circular polarization, it will be received without loss in one output of the receiver, whereas if it radiates any linear polarization, the loss will be 3 dB in either receiver. An attractive possibility for a beacon is to radiate LHC polarization, say, to represent a zero, and RHC polarization to represent a one. In this way the sender, realizing we have thought about it too, is assured of a binary channel that does not degrade the signal detectability when modulation is present.

Since such modulation may well be present, the pattern detector must examine the output of two MCSAs, one for the LHC receiver, the other for the RHC receiver in searching for regular pulse trains. The pulse trains in either polarization would have as many as half of the pulses missing in an irregular way; only the combined output is regular (see Figure 8).

Since the noise in the LHC and RHC receiver is independent, power spectra out of the two MCSAs may also be added sample by sample. The SNR improvement would then be nearly 3 dB at high SNRs and 1.5 dB at low SNRs. Thus, linearly polarized pulses and CW signals could be received with a maximum of 1.5 dB effec-



It takes only 1.2 ms for the system computer executing the present pulse-train detection algorithm to find and mark the regular train of five pulses shown above. To do so for 1000 samples, over 10 MHz of spectrum resolved to 1-Hz bins would take 38 hours. To do the job in real time (1000 s) will

require a special-purpose highly concurrent computer. Notice that the pulses in the train differ in power because the noise amplitude has added to the pulse amplitude in some cases and subtracted in others.

FIGURE 8. Pulse Train Ahoy!


tive polarization loss. The cost of the improvement is, of course, a 50 percent increase in the data processing required in the pattern detector.

SETI SYSTEM OPERATION

The entire receiving system consisting of the antenna, the receiver, the MCSAs, and, in the case of the targeted search, the pattern detector is under the control of a system computer, so that the operation is essentially automated. This is important, since negative re-

sults are to be expected for months, or years, on end. Human operators would become frustrated and lose their diligence, but a computer will not (see Figure 9, p. 1162).

In the sky-survey mode, the computer scans the antenna in a prescribed pattern and notes the coordinates and strength of any above-threshold signal. In the targeted search, the coordinates of the target stars are read from a tape file in a sequence that minimizes slewing. If a suspicious signal is found, the search is halted until



Signals from an existing radio astronomy antenna are amplified by a low-noise dual-polarization receiver. The two outputs, representing radiation received as left and right circular polarization, are delivered to MCSAs that resolve several megahertz of spectrum into several million channels as narrow as one Hz. Patterns appearing in successive spectra due

to CW signals or regular pulse trains (possibly drifting in frequency) are detected automatically. If the probability of their being due to noise is low enough, the system computer initiates a series of tests to prove or disprove their extra-terrestrial origin. If no signals are found, the computer selects a new target and continues the search.

FIGURE 9. An Automated SETI System


verification tests have been made. These tests, which are similar for both search modes, include

- consulting files of known radio-frequency interference,
- determining that the signal disappears with the antenna off target,
- executing a small-angle conical scan to determine that the signal is on main antenna lobe,
- determining that the sidereal position of source is fixed, and
- determining that no known spacecraft lies in that direction.

During this test period, the signal is recorded so that possible modulation can be found later. If all tests are passed, the computer asks for human help. Assuming the human help confirms the apparent validity of the detection, the first check would be to inquire if other observatories also find the signal, given its coordinates, frequency, and other pertinent characteristics. This would rule out local interference and, more likely, fraud.

NOTE ADDED AT PRESS TIME

The 74,000-channel prototype of the MCSA, described in the text, is now being field-tested at Goldstone, California. Pioneer 10, now outside the solar system, is still radiating its 1-W carrier from an antenna with 32.6 dB gain. The signal, received on a 26-m antenna, is too weak for the station to achieve lock. However, using a stable frequency synthesizer as the local oscillator, the signal was clearly visible on the MCSA display terminal—see Figure 10. The top window shows noise power in 144 bands each 1152/2 Hz wide. The center window shows power in 1024 bins of one band, each bin 1/2 Hz wide. The lower window displays power (as brightness, or pixel area) in 206 bins in the grey band of the center window. One hundred successive spectra are shown. The figures at the bottom are a menu of icons selected by cursor and used to modify the display.



Pioneer 10 radiating about 1 W at a distance of 5.3 billion km (3.3 billion mi.) produced the above trace on the MCSA display. The frequency scale is reversed; lower frequencies are to the right. The MCSA was clocked down to $\frac{1}{2}$ normal speed to give $\frac{1}{2}$ Hz bins and a sample every 2 s. The

downward drift of about 0.2 Hz/s in the received frequency is caused mostly by earth's diurnal Doppler. (Signal-to-noise power ratio here is 6 dB or more, so trace is more visible than those in the computer-generated Figure 4.)

FIGURE 10. Pioneer 10 Received on the MCSA

FURTHER READINGS

- Billingham, J., Ed. *Life in the Universe*. MIT Press, Cambridge, Mass., 1981.
 NASA. Project Cyclops. Publ. CR114445, NASA, Moffett Field, Calif., 1973.
 NASA. The search for extra-terrestrial intelligence: SETI. Publ. SP 419, NASA, Moffett Field, Calif., 1977.
 NASA. SETI science working group report. Tech. Pap. 2244, NASA, Moffett Field, Calif., 1983.

CR Categories and Subject Descriptors: D.4.1 [Operating Systems]: Process Management—concurrency; I.5.4 [Pattern Recognition]: Applications—signal processing; J.2 [Physical Sciences and Engineering]: astronomy

General Terms: Design

Additional Key Words and Phrases: digital Fourier transforms, finite impulse-response filters, interstellar communication, Search for Extra-terrestrial Intelligence, signal detection, spectrum analysis

Authors' Present Address: D.K. Cullers, Ivan R. Linscott, and Bernard M. Oliver, NASA Ames Research Center, Mail Stop 229-8, Moffett Field, CA 94035.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

Copyright of Communications of the ACM is the property of Association for Computing Machinery and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.