PROJECT SENTINEL: ULTRA-NARROWBAND SETI AT HARVARD/SMITHSONIAN

Paul Horowitz
Department of Physics
Harvard University
Cambridge, MA 02138

John Forster
P.O. Box 48
MIT Branch Post Office
Cambridge, MA 02139

ABSTRACT. We have been searching the northern sky for ultra-narrowband carriers of intelligent extraterrestrial origin at 21cm wavelength. The dual-polarization 65,536-channel receiver has a resolution bandwidth of 0.03Hz, matched to the Drake-Helou spreading of the interstellar medium. This maximizes signal/noise ratio, and simultaneously rejects carriers of terrestrial origin. The results of 15 months of observation are summarized.

This is a literal transcription of the talk, slightly edited for readability.

I am going to tell you about project Sentinel, which is the code name for the SETI that we are doing at the Harvard University/Smithsonian Astrophysical Observatory 84-foot dish out in Harvard, Massachusetts. This is the same site that we'll be visiting Thursday afternoon. This research is supported by the Planetary Society, and we began our continuous observations a little over a year ago, in March, 1983. I would like to say at this point that Bob Dixon at Ohio State University has been doing a continuous all-sky 21-centimeter search for 10 years, to correct the impression in the program booklet that we are unique in doing a round-the-clock search.

The basic idea of our search is, as Jill Tarter pointed out, to do extremely high resolution multichannel spectral analysis for the detection of pure carriers -- CW signals -- with the best signal/noise ratio we can achieve with the dish we have. This turns out to have an incidental benefit of giving us very good rejection of radiofrequency interference, for reasons I'll explain shortly. To get this maximum signal/noise ratio we matched the resolution of our multichannel receiver to the bandwidth of the interstellar galactic medium; let me explain what I mean by that: This is the Drake-Helou work of 1976 that Phil Morrison talked about last night. Here (figure 1) is a schematic representation of the mechanism by which a pure carrier gets broadened by passage through the interstellar medium. They call the process "phase modulation broadening through multiple scattering in a turbulent inhomogeneous

M. D. Papagiannis (ed.), The Search for Extraterrestrial Life: Recent Developments, 291–303.

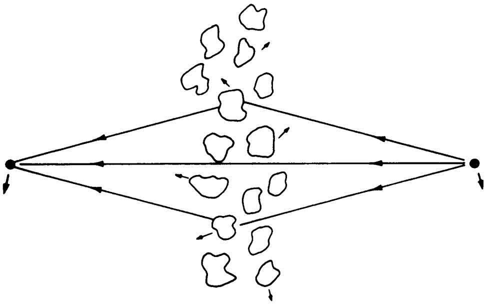

Figure 1. Phase modulation broadening of a carrier by multipath scattering in a turbulent inhomogeneous medium.

medium" (I guess we should make an acronym for that!). Anyway, the basic idea is that if you have a point source of radio waves you could have a direct path, as indicated, but you could also have a scattered path if there is ionized stuff around. If that stuff is moving, then you'll get doppler shifts from each scattered path and the sum of all those paths will give you an effective broadening, a stochastic sort of broadening. Drake and Helou calculated this effect because it seems interesting for the problem of interstellar communication. These are, by the way, the same moving blobs of stuff that also give rise to interstellar scintillation, as we see with radio signals from pulsars. Here's the graph out of their paper (figure 2) showing the net spreading in frequency as a function of distance (assuming average sort of parameters for the amount of this ionized stuff, based on pulsar observations). If we're talking a distance that's out to perhaps a kiloparsec,

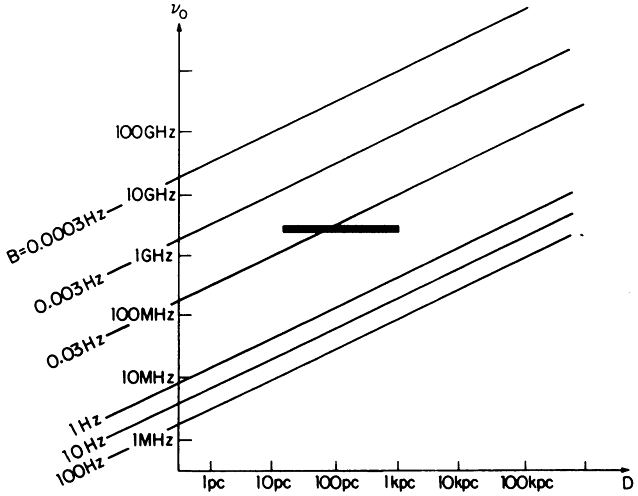

Figure 2. Spreading of a radiofrequency carrier by propagation through the galactic medium (After Drake & Helou, NAIC report 76, 1976).

and if we're talking an observing frequency of 1.4 Gigahertz (that seems to be the favorite today), then we're talking spreading bandwidths of order 0.01 to 0.1 hertz. That's a rather small spread, it certainly is tiny on the scale of the kind of resolution bandwidth that you talk about in radioastronomy, but nevertheless, that is the ultimate limit. You can't propagate a pure carrier and have it come out narrower than that as long as it has to propagate through this stuff.

What we would like to build, then, in order to optimize detected signal/noise ratio, is a spectrometer with resolution of order of tenths or hundredths of hertz. Now the problem here is that we also would like to be able to cover a significant portion of the radio spectrum because we don't know where these guys are transmitting. A resolution of a hundredth of a hertz or a tenth of a hertz makes it extremely difficult to do even the whole "water hole" of 1.4 to 1.7GHz. With 0.03 hertz, you're talking a 10 *giga*channel analyzer to cover the water hole instantaneously. But this whole idea that if you want to go for optimum signal to noise ratio you should match your spectrometer resolution to the properties of the interstellar medium, if that's really what the extraterrestrials do then it makes even more essential the idea of "magic frequencies". In other words, magic frequencies and ultra-narrow bandwidth reinforce each other, because then we'd get away from needing a 10 gigachannel analyzer.

Now, even with the assumption of magic frequencies, it turns out that the doppler shift uncertainties caused by relative motions are large compared with attainable bandwidths, if we insist on spectral resolutions comparable with the Drake-Helou spreading. For instance, at 1420MHz a kilometer per second is a half a million channels at 0.01Hz resolution. The earth's orbital velocity turns out to be 15 million channels. So if you don't somehow compensate for these dopplers, you're going to need multi megachannel analyzers just to have the requisite coverage. What do we do about this? Tomorrow I'll describe the system we're now constructing that does an end run around this, along the line suggested by Phil last night, but for now I'll simply describe what we're doing with the current system, namely we correct for the doppler shifts caused by our motion relative to the line of sight. We know the direction in which our antenna is pointed, and we know the earth's ephemeris with great accuracy, and so we can correct our receiver easily in real time, for, let's say, our motion relative to the heliocenter. What about the guys at the other end? Well, if they are transmitting in a particular direction with a beamed antenna, they can certainly do the same thing, and that would seem to be a reasonable thing to do: Why not simplify the task by avoiding signals with several FM motions due to these rotational shifts superposed? So let's assume they take out their site doppler as we take out ours. That still leaves one last term which is the motion of the stars with respect to each other, and you know that random peculiar velocities along the line of sight of nearby stars is something like 20km per second; so we're talking another 100 kHz or so of frequency uncertainty, and therefore it would certainly be good if we can get rid of this term. Well, there's two things you can do here: You can look at each star measure its radial velocity through spectroscopy, and compensate for the corresponding radiofrequency doppler. In our current search we don't do that, we say instead, "well look, they're smarter, they're doing the transmission, they're older, all the good arguments about the asymmetry of SETI at this stage in our technology, let's let them compensate the stellar radial velocity".

Now, you'll realize immediately that this puts an additional constraint on the channel, mainly we must assume that they are transmitting not just in this direction, but at this star specifically, so they can make that radial velocity compensation. So everything I'm saying from now on assumes a search for civilizations that are targeting our particular star. There are good reasons why a civilization might want to do that: If it is nearby it may have heard our leakage radiation; or it may know for other reasons that we're technological, or may be pretty soon; or it's possible that proselytizing of the "galactic club", if such exists, may happen on a chapter by chapter basis, always done from nearby. In such cases, a compensated beamed beacon is not the dumbest thing in the world to do.

Well OK, let me talk now about the doppler corrections that we have to do and this rather interesting fact that Jill Tarter alluded to about the rejection of interference. The earth's orbital velocity turns out to be almost precisely 10^{-4} of the speed of light. And that corresponds to ±150kHz of frequency offset at 1.5GHz. That would require 15 million channels of 0.01Hz if we didn't correct it, but we do. The earth's spin produces a maximum doppler shift of just 2.5kHz, and again we compensate our receiver frequency. But it turns out that this spin term has the larger rate of change of observed frequency with time -- a doppler chirp because it's a change in the doppler shift with time. The time rate of change of the received frequency is simply $df/dt = (f/c)(dV_r/dt)$, where V_r is the radial velocity. This is maximum (although the doppler shift itself is zero) when the source is observed overhead at the equator. If you put in numbers for 1420MHz, you find that the rate of change of a received signal overhead at the equator is 0.16Hz per second. And, of course, if you're anywhere else you have to put in these generalizing factors of cosine of this times cosine that, whatever this and that are called. I guess that one is called declination (or maybe it's the complement of it. I don't know, I'm not an astronomer actually). And that one there is called lattitude (or maybe its the complement). Anyway, you know what those things are. Cosine of this times cosine of that, but cosines are always near one and it really doesn't pertain to the argument very much. And the other curious fact is that observed frequencies are always going down; I'll let you think about that one because it seems paradoxical at first. How do they ever get back up to where they started, you can answer that one with a moment's thought.

Anyway, here's the business about the RFI rejection: Let's say we're talking 0.01 hertz resolution bins, because we want to match the interstellar spreading as instructed by Drs. Drake and Helou. That means we have to take a hundred seconds' worth of data, which at this chirp rate produces 16 hertz of drift of a celestial signal. Well 16Hz is 1600 resolution bins, so a signal transmitted from "out there" at fixed frequency with respect to our heliocenter, say, arrives here sweeping through 1600 bins in one or two minutes' time. Obviously, something has to be done about that -- that's the bad news. The good news is that if we do correct for that we will see a celestial source satisfying all these requirements as a pure signal but local interference will now be swept out by our moving receiver. To put it another way, we're looking for a signature of a signal coming from space -- a rate of change of observed frequency because of the earth's spin.

This idea was tested out first in 1978 with a software off-line multi-channel spectrum analyzer of 64,000 channels at 0.016Hz resolution. This was a search I did at Arecibo at the invitation of Frank Drake when I had a sabbatical to squander. The receiver was a standard low-noise heterodyne front-end, with quadrature baseband signals sampled at 1kHz and written onto 9-track tape. We couldn't keep up with the data in real time, but we did it off-line. And it gave us a chance to try out some of these ideas. In one test we generated a weak chirping signal to match our chirping receiver and put it down 30dB relative to power in the pass band of 1kHz; It showed up clearly as an 11 sigma peak, slightly spread into 3 channels (of 64K), at the correct beat frequency. So chirping signals do indeed ring the bell if you beat against the chirping receiver. Just as long as the receiver is well behaved and the receiver's local oscillator can sweep smoothly that should of course be the case. In a second test, we put in a fixed frequency signal at 1000 times the power level (0 dB relative to power in the entire 1kHz band). It's not at a fixed beat frequency of course, because it now chirps with respect to the receiver, and, lo and behold, it doesn't show up anywhere. The strongest spectral feature in that test was a noise signal at 5.4 sigma at a completely unrelated frequency. The conclusion seems to be that chirped receivers, which you need for observations at spectral resolutions comparable to the Drake-Helou value of 0.01Hz, are highly insensitive to terrestrially generated interference. By the way, here's what a signal looks like in the frequency domain (figure 3). This is a pure carrier inserted at -20dB level,

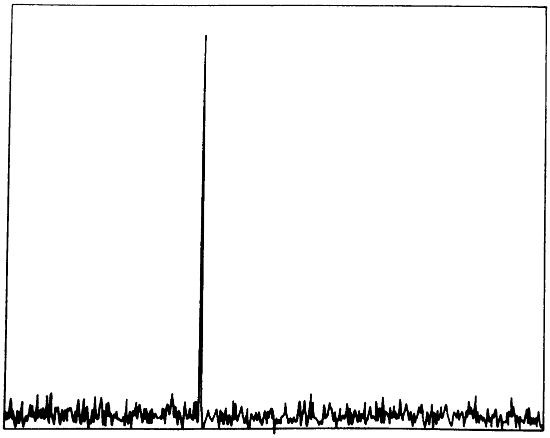

Figure 3. Portion of a 64K-point spectrum of a carrier injected at -20dB. 512 channels are shown.

processed with the 64K transform, again at Arecibo. As usual, in SETI, the only signals you find are ones you make yourself. This graph is a little piece out of a 100-foot long graph if you plotted the whole thing. But it shows that even at 20dB below noise, it looks pretty good, as it should with 64,000 channels. Oh, by the way, while we were at Arecibo, we looked at the 200 nearest F, G and K dwarfs to the sun for the "real thing", and we didn't find anything.

The current project was actually begun at NASA-Ames and Stanford University in 1981-82 when I had a sabbatical there, and we called it "Suitcase SETI". This is a good time to mention the participants in the project: The Planetary Society paid for our equipment and operation, the observatory belongs to Harvard University and SAO, a NASA fellowship supported me in California and NASA paid for about 60% of the original receiver, and Barney Oliver at Hewlett-Packard helped us out with an oscillator and synthesizer. From NASA, Stanford, and Berkeley we were helped by Peter Backus, Kok Chen, Ivan Linscott, Tap Lum, Alan Peterson, and Cal Teague. At Harvard we've had the expert help of Dave Brainard and John Forster. And at the Oak Ridge Observatory, Arnie Aho is responsible for the nice tiled floor you'll find in the lavatory out there, his wife did the lace curtains, Gene Mallove helped revive the sleeping beast, Mal Jones (who is sitting in the second row here) has painted us pretty silver and white and purple splotches on our telescope, testing how to paint over rust, John Ball (the last legitimate user of the dish) showed us how to make it go -- and Skip Schwartz and Dick McCrosky keep it going.

Suitcase SETI was simply a real-time hardware version of the Arecibo experiment. The trouble with off-line spectral analysis is that you get way behind in analyzing your data, tapes pile up in the hallway, and you end up with lots of data you never analyze, along with possible detections that it's too late to follow up. At Stanford we built a dual 64,000 channel hardware receiver with agile local oscillator, that is, a receiver able to track the signature of an external signal. We put in on-line signal recognition algorithms and archiving of all the data on videotape, and so on. The entire system itself is contained in three suitcase-sized boxes, as shown in figure 4.

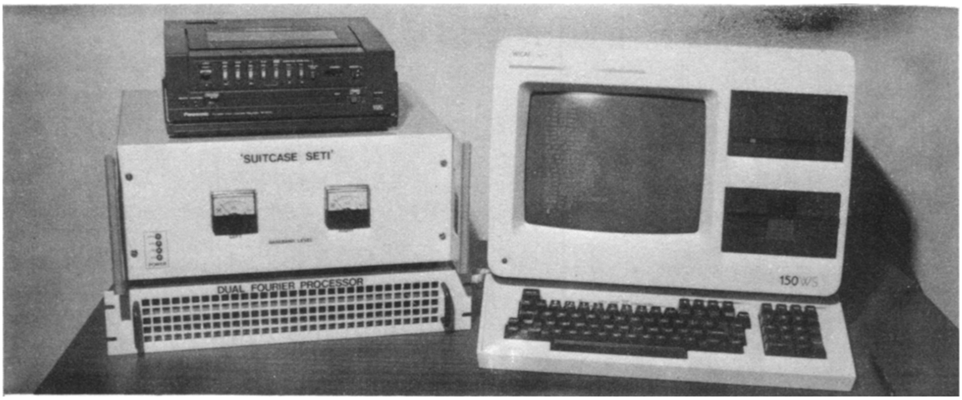

Figure 4. "Suitcase SETI": A portable multichannel receiver with real-time signal recognition and archiving.

There it is, that's Suitcase SETI, kind of a big suitcase, but if you have a few mules, you can do OK. Most of the RF hardware and the swept receivers are in the box at the left. The dual Fourier processor is just what it says, and the commercial computer on the right looks at all the data, controls the oscillator, runs the video tape archive, draws graphs, signals and beeps at you when it finds something.

This system had its maiden voyage at Arecibo in the spring of 1982 and we again tested out the hardware, looked at 250 nearby stars at hydrogen (1420MHz), and at twice hydrogen (2840MHz). We didn't find any signals, but we didn't find any false alarms either, as was true also of the earlier search at Arecibo. It really does show that this rejection of terrestrial interference is genuine. In fact, you really have to put in some signals to make sure the system is working because it seems to be so clean in terms of interference. We did look at one radiofrequency source, W49OH, one of the OH masers, which is known to have some nice sharp features by ordinary radio astronomical standards. We took a whole set of contiguous spectra, each one having 64,000 points and typically showing just a piece of a curve, and altogether we created a data set which, if you plot the spectrum of just the 1665 line at 200 points per inch (high resolution Versatech plotter), the graph would stretch across the 1,000 foot dish! We were disappointed in not seeing any hundredth of a hertz features, but of course, if they had been there, someone would have noticed before and turned up the resolution.

That's basically where the Sentinel Project was coming from, mainly a piece of hardware that knows how to look at 2 x 64,000 channels and what to do with it, now that we've used up our annual allocation of Arecibo telescope time. What we would really like to do with it is look at a million stars, or ten million stars. Two hundred stars is fun, but it's probably not going to be enough. No matter how optimistic you are about the Drake equation, you are in trouble with 200 stars. Well how do you look at a million stars? It turns out, and Mike Davis pointed this out to me, he said "you foolish boy!" (or words to this effect), "with an 85-foot telescope there aren't a million resolvable points in the sky, why don't you just do a meridian transit scan of the whole sky?" So, although we sometimes like to say that we are looking at a million stars, we're really doing an all sky survey, and Barney Oliver refers to this as a "directed all-sky search".

Anyway, with funding from the Planetary Society we fixed up the Oak Ridge 84-foot telescope, built receivers, put a new roof on, and began a full-time search. We do a search of the entire northern sky (or about 80% of the entire sky) in about half a year, covering one beam width, that is, a half-degree circle around the sky, each day. Our search began in March, 1983, doing a transit search at 1.4GHz (the dish is good to 5GHz). Our Berkeley L-band "lumplifiers" have a noise temperature of 50^{o}K (uncooled), giving us a system temperature something a little under 100 degrees. We do a pair of 64K transforms of resolution 0.03Hz, covering 2kHz in each antenna polarization. Since we are looking for predopplered transmissions at magic frequencies, we chirp the receiver to compensate for our site motion. A data run at 0.03Hz takes 35 seconds, and a source, if there is such a thing as a source, is in the beam for three minutes (because the beam is half a degree).

The system searches for large peaks, and archives anything suspicious, along with observatory parameters.

Let me just show you the obligatory block diagram (figure 5) and then I'll show you the kind of data we have. Starting at the dish,

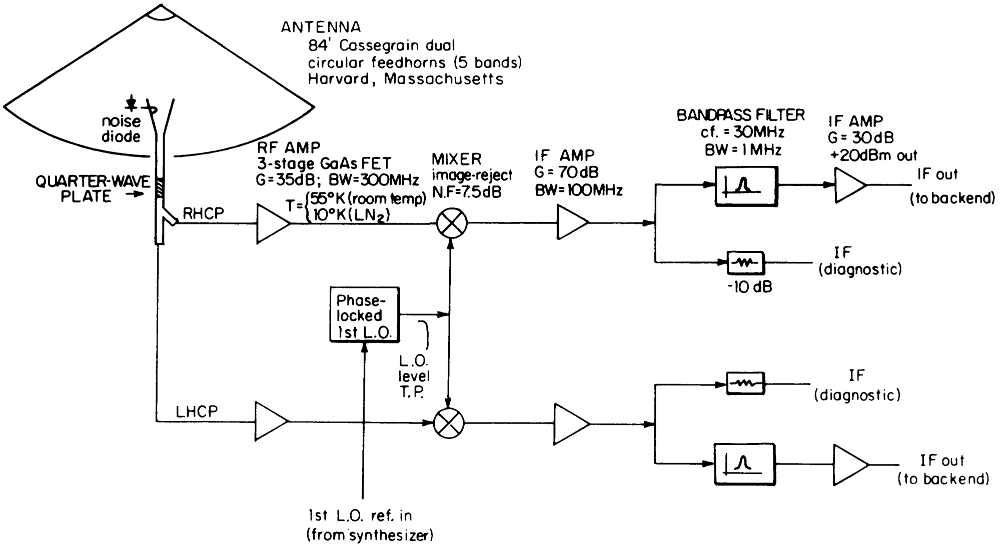

Figure 5. RF frontend of Sentinel.

the signal is split into two channels, according to polarizations. Then it is amplified by a pair of GaAs FET low-noise amplifiers, and mixed with a phase-locked LO to 30MHz IF, where it is amplified and filtered further, then sent through hardline coax to the control room. So far, it is a straightforward single-conversion heterodyne system. The receiver backend (figure 6) mixes the 30MHz IF to baseband in quadrature mixers driven by a computer-controlled frequency-agile 2nd LO. The I and Q baseband signals then pass through the usual anti-aliasing filters, sample/holds, and analog/digital converters, finally feeding the digital FFT processors with 8-bit sampled voltages at a 2 kHz rate. Then off you go to the central computer with your 64K-point spectra.

Woody Sullivan asked me to please explain how we look for peaks, what is our algorithm, is it any good at all. Our algorithm is really very simple since we don't seem to have much interference. Simply look for large peaks. So we calculate a moving baseline on the 64,000 frequency points, keeping a baseline that follows long wavelength features, and we simply look for high channels. If we see a big one compared with the fluctuation from the background, we flag it and write down some summary information about the run, including the largest peaks, their frequencies, where we were pointed, and so on.

Let me show two photographs. The first one (figure 7) is a photograph of another Harvard search which hasn't been shown yet at this meeting, so I can't resist. This is the first search for 21cm radiation at Harvard, as Ed Purcell in the front row will recognize. That's his student Doc Ewen sitting up next to the horn on the fourth floor of

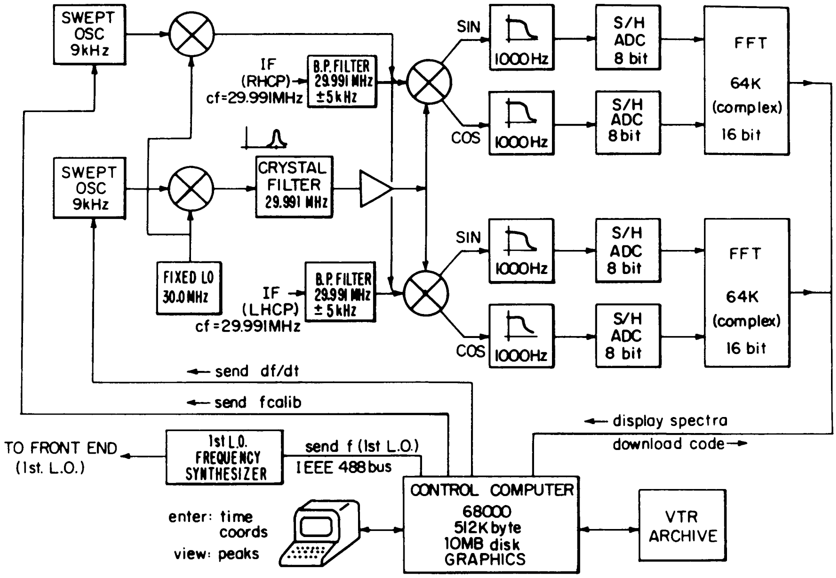

Figure 6. IF and backend of Sentinel

Figure 7. Harold ("Doc") Ewen and the historic Harvard antenna that first detected galactic 21cm radiation.

Lyman Laboratory in Cambridge, across the river. The second slide (figure 8) shows our 84 foot dish compared with a useful yard stick, my six-year-old who is slightly small for his age, but not a total dwarf, and therefore this is a realistic view of the 84 foot. Inside the control building we have the usual colorful racks of electronic equipment, computers with animated displays, etc. In our system the control computer displays a new set of data every 35 seconds, consisting of a pair of spectral graphs (broadband, narrowband) for each polarization, and a summary of the ten largest peaks -- their size and frequency. Informa-

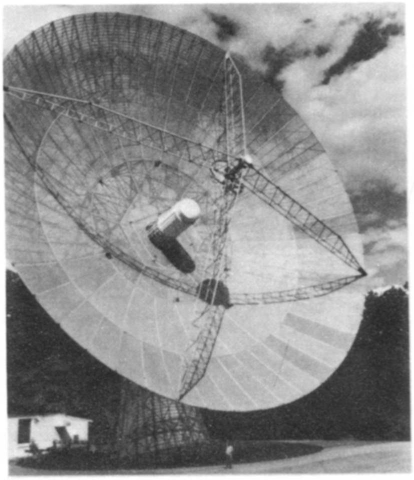

Figure 8. The Harvard/Smithsonian 26m dish with a convenient size scale (Jacob Horowitz, age 6).

tion about integrations containing unusually large peaks is written onto disk, so we can call it up during our periodic visits to the station.

Let me now summarize by telling you the results of 15 months of running, I think there are some interesting lessons here about SETI. The first comment concerns the price. Contrary to popular belief it hasn't been as expensive as Abbett's cartoon in the Boston Herald (figure 9) would suggest; in fact we have been running within our budget, which is something like $20,000 a year for keeping this kind of equipment going and making repairs, keeping the building going, paying electricity. What about the sky coverage and false alarms and down time? Here's a graph (figure 10) I made summarizing a year's worth of running. Here is plotted time of year, starting in March, 1983, when we first turned on. The bar near bottom tells you whether you're running or not;

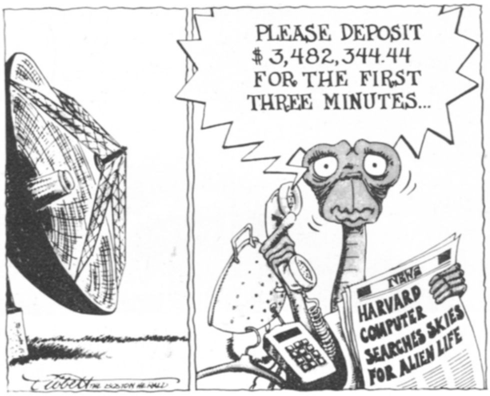

Figure 9. A popular misconception about the economics of interstellar communication (reproduced with permission of the Boston Herald).

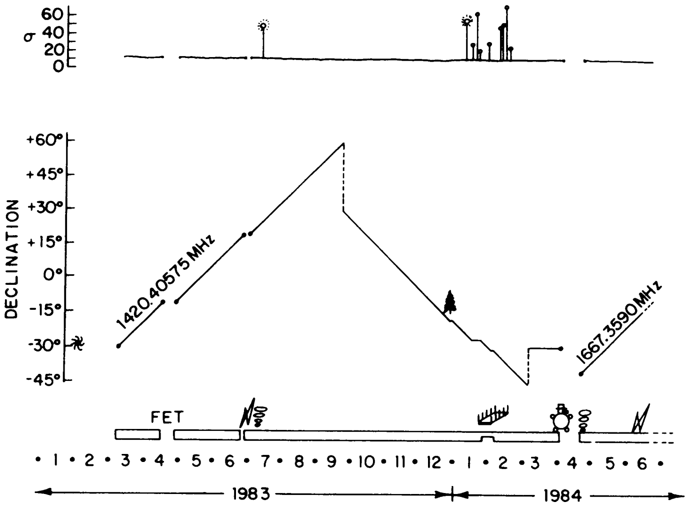

Figure 10. Sentinel: The first 15 months.

we had an infant mortality in our GaAs FET amplifiers after about 3 or 4 weeks, we could actually see the problem coming and that put us off the air for about 2 weeks. That's been about the longest down time. Things ran well until the summer time when a lightning bolt came and blew bejesus out of everything; you can see the lightning bolt on the figure. That wouldn't have killed too much, except that when the electrical company came to put a new transformer on the pole, they wired it wrong and sent 180 volts into our wall outlets, generating the smoke shown on the graph. That knocked out an extra day or two. Otherwise things have gone pretty well; in January we had a dead chip (see graph) and in March we had an incredible snow storm, the really sticky stuff that knocked down wires and again knocked out the electricity. And guess what the electric company did again, 180 volts!

The only thing I'm slightly proud of is that just a few weeks ago we had an incredible lightning storm, and it was described by Dick McCrosky, who runs the station there, in the following way: You know how you're supposed to time from the flash to the sound? He said that in this one the sound came before the flash, and he said that there were sparks jumping out of the outlets 3 or 4 inches. Well, it didn't knock us out this time, because we put a bunch of extra protective do-dads in, which I'll be happy to describe to anyone who's interested. Overall, the system has been extremely reliable.

I've also plotted sky coverage here, with the galactic center indicated by that cute little crab-like object there. We started at 1420MHz at -30° declination and began moving north. Every time we have an outage, of course, we stopped advancing declination, as indicated. But we basically covered up to +60 declination in a period of about 5 or 6 months. We then went down to +30° and decided to cover the galactic center a little more carefully and down to somewhat more negative declinations, so we went down to -45°. Then we went back up to the galactic center while we waited for some components for the OH receiver to arrive. Now what other things happened here? Well, at the top I've plotted "signals", the biggest daily signal in sigma, typically somewhere around 10 sigma for a day; note that these are not Gaussian statistics (because of squaring and summing), so don't be worried about a 10 sigma peak, it's just a noise tail. But here on July 6 or 7th we had a 49 sigma peak, and a similar event again in January. And what is that? Well it turns out we're at declination 22° and the sun had just finished peaking at 23 1/2° and heading back down, and it crossed our beam that day -- it had to happen sooner or later -- and so we saw the sun. I'll show you later why we know it's the sun and what its signature looks like. Anyway the sun passed us twice. And here in January and February, a different group of false alarms, and I'll show you next what they look like in detail, and why they cannot be extra terrestrials Luckily they turned off, because we were going to have to go out there with a shotgun to find out who it was. Apart from these two kinds of events, we've had essentially no false alarms and very few breakdowns.

In the last slide (figure 11) you can see the signature of several kinds of signals, which again I have to emphasize have nothing to do with ETI, but do show some interesting lessons. First of all, for calculation, the top graph shows what a drift scan through a radio source, Taurus A, looks like. This should give you a feeling for the intensity profile of the antenna, and I've plotted actual measured intensities from a drift scan, but only at the intervals that we normally have spectra out of our SETI system. So, if in fact we had a CW signal coming from an extra terrestrial, it ought to have this amplitude profile as it passes through the antenna beam. The middle graph shows the sun,

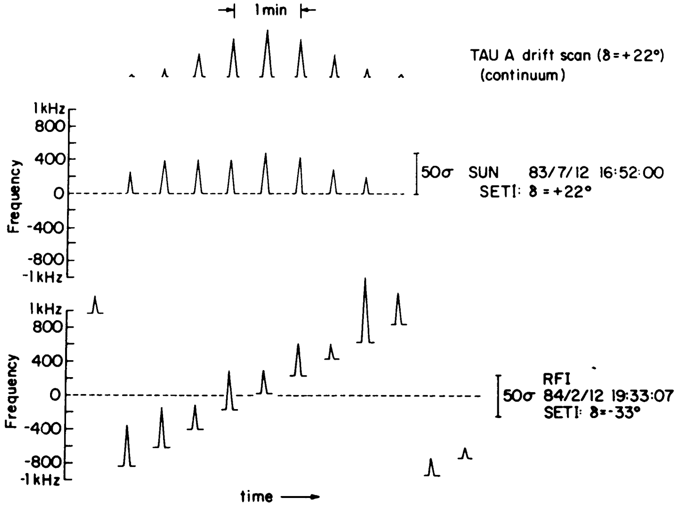

Figure 11. Frequency/intensity drift-scan signatures of three kinds of signals.

as processed by the SETI detector; it has a plausible amplitude envelope that is maximum at the center. But look at the frequency, it's always exactly at 0 Hz; that's of course suspicious. What it turns out happened here is that the signal from the sun is so intense it's simply saturating the IF amplifiers, desensitizing them and letting the small offsets in the amplifiers come through as DC. And of course the declination is right, and so is the transit time, which convinces us it's really the sun.

The bottom graph shows the "false alarms" from Jan-Feb, 1984. This is most interesting, they are marching through in frequency! Because we look at a slightly different section of sky for each run, we shift our center frequency for each integration so that a source in the center of our antenna beam, transmitting helocentric 1420MHz, will be in the middle of our backend passband. As a result, if there were a true cw radio source transiting through the beam profile, it would appear to step in frequency in each successive integration. By the way, notice the aliased frequencies at the extremes. The signal shown in the bottom graph has the stepped frequency behavior we expect, but it doesn't have much of an amplitude envelope; it seems more or less level to within statistics here. However, the real proof that this signal is terrestrial comes from its observed frequency and direction: It turns out that this signal comes from eight different places in the sky and yet it was always received at the identical site frequency. That would be a pretty strange sort of acquisition beacon for an extra terrestrial civilization to send! So it's apparently local interference and somehow getting through our otherwise invincible chirping receiver.

Well, let me just end by summarizing our Sentinel search. Remember that this search is very restrictive in that we're only really sensitive to magic frequencies, choosing one or two favorites per year. We require continuous carriers rather than pulses or chirps, and the carriers have to be beamed at us, for the reason that the senders have to remove the doppler shift caused by radial motion between heliocenters. Tomorrow, John Forster and I will be talking about our new 8 million channel system, which gets around this restriction. The concentration on narrow carriers gives us rather impressive sensitivity: It turns out a tenth nanowatt incident on the entire earth ($\approx 10^{-25}$ W/M^2) in our band pass would trigger a 20 sigma event. Using reasonable figures of merit, one can restate this sensitivity as 10 kilo Ozma/minute (where are you, Frank Drake?) And the coverage is approximately 80% of the sky, in half a year, with frequency coverage (in dual polarizations) of 2kHz at a resolution bandwidth matched to the Drake-Helou bandwidth of 0.03Hz, or, equivalently, 128,000 channels. Tomorrow we will describe our new system, which lifts several of the more oppressive restrictions. Thank you very much.