Strategies, Implementation and Results of BETA

Darren Leigh

Mitsubishi Electric Research Laboratories, Cambridge, MA 02139, USA

Paul Horowitz

Physics Department, Harvard University, Cambridge, MA 02138, USA

Abstract. The Harvard University / Smithsonian / Planetary Society *BETA* project is an all-sky, narrow-band, microwave search for extraterrestrial intelligent signals. It has been operating nearly continuously for the last four years during which time it has automatically scanned the sky visible from Agassiz Station (+60° to −30°) over the entire waterhole (1400–1720 MHz) four times. We discuss *BETA*'s search strategies, our implementation and the results of how these fared in the observatory's interference environment. We also present qualified limits on the prevalence of transmitting civilizations given our (current) negative results.

1. Introduction

The Billion-channel ExtraTerrestrial Assay (*BETA*) project is the first all-sky, all-waterhole search for microwave signals from intelligent, extraterrestrial sources. It was developed in the early 1990's as a replacement and upgrade for our previous experiment, the Mega-channel Extraterrestrial Assay (*META*). *BETA*'s design was greatly influenced by our experiences with *META*. Our goal was to create a search system with much wider frequency coverage than *META* had, but which was less prone to interference and unexplained "one-shot" events. The *BETA* architecture and implementation have been described in (Leigh and Horowitz, 1996) and are discussed in depth in (Leigh, 1998). In a nutshell, the system comprises 250 million 0.5-Hz channels, divided into three antenna beams; the resultant 40 MHz instantaneous bandwidth covers the microwave "waterhole" in eight hops of two seconds each. This 16-second cycle repeats eight times during the sidereal drift of a point source through the antenna's 0.5° beams. This document concentrates on some of *BETA*'s novel features and how well they performed.

2. Interesting Architectural Features and Their Results

Building on past experience, we designed *BETA* to be more robust than our previous efforts. It is set up to provide sufficient information to either prove or disprove the source of a signal. It allows immediate, automatic follow-up so that

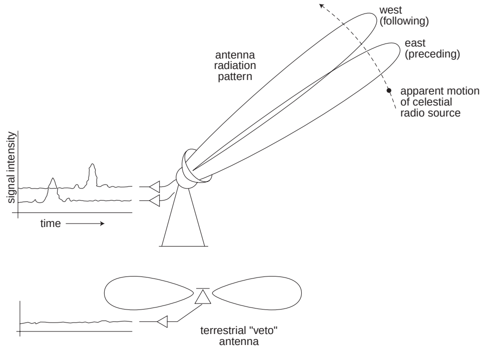

Figure 1. *BETA*'s three beam feed system

short-lived signals do not remain unproven mysteries. Three of the architectural features used to provide this capability are described below.

2.1. Three beam feed system

An extraterrestrial signal has one important characteristic that terrestrial RFI never has: it will always appear to be coming from a sidereally fixed point on the sky, unaffected by the position or motion of the earth. We designed *BETA* to use this characteristic as its main RFI filtering strategy. *BETA*'s feed system has three beams. Two adjacent beams point at the sky: one slightly east and the other slightly west. Any signal coming from a fixed sidereal position will trace out the characteristic pattern of the beam lobes, first in the east beam and then in the west.

A third low-gain, azimuthally omnidirectional antenna looks predominately at the horizon to pick up any interfering terrestrial signals. This provides a fast way to veto strong interference, in hardware, early in the recognition process. This works well for strong carriers, but is not perfect. Some satellite signals may appear in the sky horns but be too weak for the terrestrial horn thresholds. Also, although RFI signal strengths are approximately the same for all of the feeds (the terrestrial feed and most of the telescope sidelobes have gains of around 0 dBi), the terrestrial feed has more thermal noise because it looks mostly at the ground, raising the noise floor and thus desensitizing it.

The spectra from all three beams (east, west and terrestrial) are computed simultaneously and synchronously so that any frequency comparisons can be quite accurate. *BETA* was designed to follow the time history of these spectra over the course of several minutes (enough time for a sidereal source to transit both sky beams) and compare them to the expected behavior of an extraterrestrial signal. If the result suggests a strong candidate, the antenna can be moved ahead in hour-angle by several beamwidths to allow the source to pass through the stationary beams again. This "leap-frogging" can be done up to six times in succession.

Overall the three beam system performed admirably. None of the spectral features we detected had the proper characteristics for a sidereally stationary source (first east, then west, never terrestrial). Thus, despite the fact that we detected many candidate signals, we were not left with any "mysteries". The three beam system does have some disadvantages, however. An ETI signal must remain above threshold for six to ten minutes in order to transit both sky beams, so the system is insensitive to short-lived signals. Also, ETI signals that happen to be at the same frequency as RFI picked up by the terrestrial feed will be eliminated.

2.2. Real-time finite state machine

Most of the signal analysis and recognition in *BETA* is implemented in software. Due to the real-time nature of the system, the thresholds have to be set reasonably high so that the software can handle all of the candidate signals. In order to do some basic signal recognition at lower threshold levels, we implemented a programmable state machine (SM) which can examine *all* of the data for elementary patterns. It determines whether a candidate is interesting enough to forward to the analysis software on the host PC.

The SM receives as input: how many thresholds were exceeded by the signals in the three beams, "time based state" (information it stored the last time a particular frequency was visited), and "frequency based state" (information it stored from the previous frequency bin). Based on this information, the SM generates new time based state, frequency based state, and a bit specifying whether the data of this frequency should be forwarded to the host PC for further software analysis.

The SM allows great flexibility in making intelligent decisions at speeds with which software cannot compete. For example, the frequency based state can be used to detect a signal that happens to lie between two bins, so that it does not exceed the highest threshold, but exceeds the second highest threshold for two consecutive frequency bins. The time based state can be used to monitor a signal that is crawling into the edge of a beam, but has not yet arrived at the center and so has not exceeded a high threshold, but has repeatedly exceeded lower thresholds. Since the SM sees the comparison results from all three horns, it can perform rejection of terrestrial interference and of signals that are simultaneously high in both the east and west beams, which may indicate RFI that has not been detected in the terrestrial antenna.

We had high hopes for the state machine. The prospect of fast, very low threshold detection was extremely appealing during the design phase. The idea was to track every frequency bin's progress in hardware, looking for the east to west characteristic of a sidereal source. In an interference-free environment (or if the terrestrial veto were better) this would have worked very well; since thermal noise is uncorrelated between spectra it is highly unlikely that such a pattern would occur from chance, even at low thresholds. The problem is that interference *is* correlated between spectra. Random interference will frequently generate any of the patterns if the thresholds are low enough. This low-level detection scheme has a further drawback, namely a state machine "hit" is not triggered until the event has already occurred (and finished); the SM cannot store any of this data so there is nothing to analyze further.

The state machine can still be used to trigger regular hits and then proceed with the slot generation and analysis. It was also very useful for debugging and running various tests on the equipment. It was, however, tricky to design and debug, so we would probably not implement something this general in a future system.

2.3. Adaptive RFI filtering

There *is* intelligent life in the universe and it does transmit signals — interference! This is (and probably always will be) the number one problem in SETI observations from earth. In our search, *BETA*'s adaptive RFI filtering scheme was designed after we had some preliminary experience with the RFI environment that plagues our observatory. Before we implemented it, the interference was so severe that the system was nearly useless; with adaptive filtering the situation is tolerable.

Rather than ruthlessly notching out large, fixed frequency bands to combat RFI, the scheme temporarily masks small bands only when continuous interference has been detected there. We take advantage of the fact that interference is correlated both in time and frequency:

- If there is interference at a certain frequency at a certain time, there will probably be interference at that frequency shortly before and after that time.
- If there is interference at a certain frequency at a certain time, there will probably be interference at nearby frequencies at that time.

We exploit these correlations by breaking the FFT spectrum into small blocks, counting the number of hits we get in each block, and following the evolution of this number over time. A counter is associated with each block. If the number of hits in a block is greater than zero, its counter is incremented by 5. If there are no hits, the counter is decremented by 1. If the counter reaches 25, then that block is declared to be "notched" and no slots are generated in its frequency range. If the counter reaches zero, then the block becomes "unnotched" and hit generation is allowed again.

The scheme consistently eliminates more than 99% of the interference while masking less than 1% of the available spectrum. In order to test the adaptive filtering properly, we incrementally lowered the thresholds while watching the hits that were produced. Every time we decreased the threshold, new RFI suddenly appeared and then gradually went away. Since thermal events appear at different frequencies in each spectrum, the non-interference, thermal background got thicker each time and did not go away. Because the counters ramp up five times as fast as they ramp down, we should in principle be able to use the scheme with thresholds so low that each 2 kHz block gets a thermal hit about every five spectra. This corresponds to a threshold of $8.5\,\sigma$, or about 800 hits per spectrum, which is far more data than the rest of the system can handle.

It would not be difficult to implement the adaptive notching scheme in hardware. This would enable the system to handle large quantities of interference without overloading any of its data paths In hindsight, we probably should have done this instead of developing the generic state machine.

3. Conclusions

The current negative results of the *BETA* search allow us to set some limits on the prevalence of transmitting civilizations, with certain qualifications. From the system's sensitivity parameters we can derive a relationship between a transmitter's EIRP and the maximum distance at which we could detect it. The possible detection distance is

$$R = \sqrt{\frac{(\mathrm{EIRP})A_r\tau}{8\pi\gamma k T_N}}$$

where A_r, τ, γ and T_N are the receiver's effective aperture, integration time, threshold level and system noise temperature, and k is Boltzman's constant. The extra factor of 2 in the denominator is due to the fact that we are receiving with linear polarization, but the signal was presumably transmitted with circular polarization. With our values of $A_r = 239$ m^2, $\tau = 1/2$ second, $\gamma = 15\,\sigma$ and $T_N = 85\,\mathrm{K}$, we obtain $R = C\sqrt{(\mathrm{EIRP})}$ with $C = 1.6 \times 10^{10}\ \mathrm{m}/\sqrt{\mathrm{W}} = 1.7 \times 10^{-6}\ \mathrm{ly}/\sqrt{\mathrm{W}}$.

Like the *META* analysis in (Horowitz and Sagan, 1993), we will consider three types of transmitting "super-civilizations", similar to the Kardashev[1] definitions (Kardashev, 1964). *Type 0* civilizations are similar to ours and have total power resources of about 10^{13} watts. *Type I* civilizations have power available equal to the solar insolation on earth: about 10^{17} watts. *Type II* civilizations harness the entire power of their star: about 10^{26} watts. We will also consider two types of beacons: isotropic ($D_t = 1$) and directed beams with sufficient gains to be received here from anywhere in the galaxy ($R < 80,000$ ly). If these civilizations (lavishly) use a significant part of their power to broadcast a SETI beacon within *BETA*'s limitations (frequency range, doppler-compensated to an inertial frame, within our declination range, not conflicting with RFI, etc.), we can set the following limits on the number N of such civilizations in our galaxy (where ϵ is the earth-incident duty cycle of transmission, i.e. the fraction of time spent transmitting in our direction):

For Type 0 civilizations, $N\epsilon < 1$ out to a distance of 5.4 ly for an isotropic beacon. $N\epsilon < 1$ for the entire galaxy ($\approx 10^{11}$ sun-like stars) for directed beacons with transmitting gain $D_t \approx 83\,\mathrm{dBi}$. An 83 dBi antenna would be about 1000 meters in diameter at $\lambda = 21$ cm, and would have a beamwidth of about one arc-minute.

For Type I civilizations, $N\epsilon < 1$ out to a distance of 540 ly ($\approx 2 \times 10^5$ sun-like stars) for an isotropic beacon. $N\epsilon < 1$ for the entire galaxy ($\approx 10^{11}$ sun-like stars) for directed beacons with transmitting gain $D_t \approx 43\,\mathrm{dBi}$. A 43 dBi gain antenna at 21 cm has a diameter of 10 meters and a beamwidth of 1.2°.

For Type II civilizations, $N\epsilon < 1$ out to a distance of 1.7×10^7 ly ($\approx 10^{11}$ sun-like stars in our galaxy and $\sim 10^{12}$ in neighboring galaxies) for an isotropic beacon. At this power level, directed beacons could be detected at cosmological distances, which is not very useful for SETI considering the timescales of biological evolution.

The *BETA* system was designed to provide sufficient information either to prove or disprove the extraterrestrial provenance of a signal. We were not disappointed in this regard – the system detected *no* signals that remain mysterious. None of the archived candidates has the characteristics that we expect of an

---

[1] Kardashev's Type I is the Type 0 here. He did not specify civilizations that use their entire planetary insolation.

extraterrestrial signal. This is due to one of the following three reasons: either our assumptions about signal characteristics were incorrect, or we were unlucky, or no signals were within our capability to receive.

3.1. The Real World Intrudes

On March 23, 1999 the Agassiz Station 26 meter radiotelescope was blown over by strong winds. Damage to the dish surface and secondary support structure was substantial. We are currently raising funds to repair the telescope. In the meantime, *BETA* is turned off.

Figure 2. Fallen radiotelescope, damaged edge and secondary support.

Acknowledgments. We received support from The Planetary Society and its members, the Bosack/Kruger Charitable Foundation, NASA, the Shulsky Foundation, and Dr. John Kraus; and contributions of equipment from Micron Technology Inc., Hewlett-Packard Company, John Fluke Manufacturing Company Inc., Advanced Micro Devices Inc. and Intel Corporation.

We also received valuable assistance from Michael Davis, John Hagen, Martin Gimersky, William “Crash” Yerazunis, David Staelin, Kevin Duesman, Joe Caruso, Robert Stefanik, Joe Zajac, Greg Galperin, Derrick Bass, Neil Hendin, Nick “Phi” Schectman, Suhail Shah, Eric Wey, Dylan Manna, James Higbie, George Marcus, Jonathan Wolff, Jonathan “Jono” Weintroub, Charles “Chip” Coldwell and Ian “Max” Avruch.

The architecture of *BETA*’s spectrometer core evolved from ideas from the Berkeley SERENDIP program. Special thanks go to Dan Wertheimer for all of his help.

References

Horowitz, P. and Sagan, C., 1993, ApJ, 415, 218–235.

Kardashev, N. S., 1964, Soviet Astron. – AJ, 8, 217.

Leigh, D., 1998, Ph.D. dissertation, Harvard University. Available on the WWW at `http://mc.harvard.edu/seti`

Leigh, D. and Horowitz, P., 1996, Proceedings of the 5th International Conference on Bioastronomy, 601–610.