Targeted and All-Sky Search for Nanosecond Optical Pulses at Harvard-Smithsonian

Paul Horowitz[a], Charles Coldwell[a], Andrew Howard[a], David Latham[b],
Robert Stefanik[b], Jonathan Wolff[c], and Joe Zajac[b]

[a]Physics Department, Harvard University, Cambridge, MA 02138 USA
[b]Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138 USA
[c]SpaceDev Inc, 13855 Stowe Dr, Poway, CA 92064 USA

ABSTRACT

We have built a system to detect nanosecond pulsed optical signals from a target list of some 10,000 Sun-like stars, and have made some 20,000 observations during its two years of operation. A beamsplitter feeds a pair of hybrid avalanche photodetectors at the focal plane of the 1.5 m Cassegrain at the Harvard/Smithsonian Oak Ridge Observatory (Agassiz Station), with a coincidence triggering measurement of pulse width and intensity at sub-nanosecond resolution. A flexible web-enabled database, combined with mercifully low background coincidence rates (∼1 event per night), makes it easy to sort through far-flung data in search of repeated events from any candidate star. An identical system will soon begin observations, synchronized with ours, at the 0.9 m Cassegrain at Princeton University. These will permit unambiguous identification of even a solitary pulse. We are planning an all-sky search for optical pulses, using a dedicated 1.8 m f/2.4 spherical glass light bucket and an array of pixelated photomultipliers deployed in a pair of matched focal planes. The sky pixels, 1.5 arcmin square, tessellate a $1\overset{\circ}{.}6 \times 0\overset{\circ}{.}2$ patch of sky in transit mode, covering the Northern sky in ∼150 clear nights. Fast custom IC electronics will monitor corresponding pixels for coincident optical pulses of nanosecond timescale, triggering storage of a digitized waveform of the light flash.

Keywords: Optical SETI, interstellar communication

1. INTRODUCTION

The merits of optical SETI (as compared to microwave SETI) have been explored at two previous SPIE conferences,[1,2] and are well documented elsewhere[3–7]; but we resist unsuccessfully the opportunity to highlight a few advantages here. First, transmitted beams from optical telescopes are far more slender than their radio counterparts owing to the high gain of optical telescopes (150 dB for the Keck Telescope versus 70 dB for Arecibo at λ=21 cm); an 80 dB advantage is nothing to sneeze at! Dispersion, which spectrally broadens radio pulses, is completely negligible at optical frequencies. The capability of radio transmitters has reached a stable maturity, while the power of optical lasers has shown an annual Moore's law doubling extending over the past 30 years. And finally, the computational power and sophistication characteristic of sensitive microwave searches today is unnecessary for optical SETI. Detection can be quite simple — a pair of fast, broadband photon counting detectors in coincidence.

In this paper we describe the philosophy, design, and implementation of such a detector, in several incarnations: A targeted search, now in its third year of continuous observations, of some 10^4 sunlike stars; a synchronized twin, soon to go online at Princeton University; and a pixelated all-sky transit search, now in construction.

2. FEASIBILITY WITH PRESENT TECHNOLOGY

Before getting into lots of details, it's worth looking at the field broadly: If we wanted to *transmit*, what could we do now, using only "Earth 2000" technology? This is a useful exercise, not only as a sanity check, but also as a vehicle to help select an optimum "system" scenario; i.e., among transmit/receive possibilities, a scheme that works well at both ends is a better bet.

Direct correspondence to horowitz@physics.harvard.edu.

2.1. Transmitters and Detectors

Let us consider a civilization, at least as technologically advanced as our own, that wishes to establish contact with its galactic neighbors. Its task would be to irradiate, with a beacon distinguishable from astrophysical phenomena and from noise, the planetary zones of the nearest N stars within some range $R_{\max}$ (comparable to the average separation between intelligent civilizations). In our region of the galaxy $N \approx 10^3$ for $R_{\max} = 100\,\text{ly}$ and $N \approx 10^6$ for $R_{\max} = 1000\,\text{ly}$.

We assume that the transmitting civilization has a catalog of target stars, their positions, proper motions and ranges with sufficient accuracy to permit aiming with an error no greater than $\sim$ 10 AU when the beam reaches the target. At a range of 1000 ly, this corresponds to a proper motion uncertainty of 33 μas/year and a positional accuracy of 33 mas. The required range accuracy depends on the star's proper motion, but not its range; for example, to target the planetary zone (say 10 AU) of a star whose transverse proper motion is 10 km/s the range uncertainty cannot exceed 5 ly. These requirements are certainly within the grasp of an advanced civilization, given that our astrometry will achieve micro-arcsecond precision in the coming decades; and in any case these accuracies are relaxed if the transmitted beam is broadened to irradiate a larger zone, at the expense of received signal strength.

To send a pulse (or more generally, a packet of information of short duration) to $N = 10^6$ stars with a single laser system, the sender would probably use an assembly of fast beam steering mirrors of relatively small size and weight, in combination with a large objective that is steered slowly. Assuming that the sending apparatus could settle to diffraction limited pointing in ~ 0.01 sec (feasible by today's engineering standards), the recipient would observe an optical pulse coming from a nearby star repeated every 10^4 seconds. (This period could be dramatically reduced by transmitting only to an intelligent subset of the targets and/or by using multiple transmitters.)

The recipient would be able actually to observe these pulses, of course, only if *a)* the received fluence per pulse corresponds to at least a handful of photons delivered to the receiving telescope aperture, and *b)* the flux of laser photons, during the pulse, exceeds the stellar background. The remarkable fact, as we'll show below, is that *with only "Earth 2000" technology we already could generate a beamed laser pulse that outshines the sun by three orders of magnitude, in white light, independent of range.* One might consider this fact the "fundamental theorem of optical SETI"!

These pulses could be detected with a reflecting telescope of modest aperture, followed by a beamsplitter and a pair of photodetectors of nanosecond or better speed. (We choose nanosecond because it is roughly the speed of PMT's, and all known backgrounds disappear at this time scale). The electronics could be a pair of pulse height discriminators driving a coincidence circuit. The telescope would track the star by the photodetector's "singles" rate while waiting for the unique coincidence signature of some tens of photons arriving in each detector within the resolving time of a nanosecond. As we will see, this signature is easily detected even in broadband visible light; i.e. no spectral filters are required.

2.2. Backgrounds

In a companion paper[8] we address the obvious next question, namely the backgrounds against which the putative pulsed optical beacon must compete.

When detecting light pulses from the neighborhood of a star, the most obvious background is light from the star itself. We circumvent this difficulty by using fast detectors ($\sim$ns speed) so that the light from the star is just a slow drumbeat of essentially single photons. A G2V star at 1000 ly ($m_V = 12$) delivers $\sim 3 \times 10^5$ photons/sec to a 1-meter telescope. The Poisson-distributed arrivals do not significantly pile up; observing a single photon is rare in a nanosecond and large pulses are exponentially suppressed. By using a beam splitter and a pair of photodetectors wired in coincidence (a scheme taught by Dan Werthimer for the alleviation of many ills) the rejection of photon pileup is further reduced; for example, the rate of 2-photon coincidences in a pair of detectors exposed to Poisson distributed arrivals is $R = (r_1^2\tau)^2(r_2^2\tau)^2\tau$, where r_1 and r_2 are the individual rates, and τ is the coincidence window. For $r_1 = r_2 = 10^5 s^{-1}$ and $\tau = 1\,\text{ns}$, the coincidence rate caused by 2-photon pileups in both detectors is 10^{-5} per second, or less than one pileup per observing night. In other words, light from the parent star itself is unimportant on a time scale of a nanosecond.

As far as we can tell, both from calculation and from some 20,000 observations conducted so far, there appear to be few astrophysical, atmospheric, or terrestrial backgrounds that produce events like the nanosecond photon pileup expected from an intentional and powerful pulsed laser beacon.

2.3. A Transmission Scheme

To give a sense of the difficulty (or relative ease) of interstellar communication by optical pulses, we calculate several useful quantities for one specific transmission scheme: a “Helios” laser* beamed 1000 ly between two 10 m Keck telescopes, each orbiting a Sun-like star. We should note that this “Earth 2000” scheme is surely modest in technological sophistication and scale for a truly advanced civilization.

The transmitted beam is slender as it emerges from transmitting telescope, $\theta_b \approx \lambda_H/D_K = 20\,\mathrm{mas}$ (6 AU at 1000 ly). Its short (3 ns) and energetic ($E_p = 4.7\,\mathrm{MJ}$) pulses arrive at the receiving telescope, unbroadened in time, as a pulse of $N_R = \pi^2 D_K^2 D_K^2 E_p 10^{-4R/5R_E}/16\lambda_H R^2 hc = 1200$ photons. If the beam is broadened to irradiate a 10 AU disk, then the number of received photons drops to ~ 500 per pulse. Here D_K is the telescope diameter, $10^{-4R/5R_E} = 0.78$ is the extinction factor ($R_E \approx 1.15\,\mathrm{kpc}$ is the distance over which the intensity of a $1\,\mu$m pulse will decrease by 2 magnitudes; note that R_E is a rapidly increasing function of wavelength), $\lambda_H = 1.047\,\mu$m is the wavelength of the transmitted photons and $R = 1000\,\mathrm{ly}$ is the distance between the telescopes. The stellar background is quite small, $\sim 3 \times 10^{-2}$ photons/ns for a G2V star (thus ≈ 0.1 photons during the 3 ns duration of the laser pulse.

The interstellar medium both scatters and absorbs these optical pulses. The effects of scattering over large distances can be quite severe. It tends to reduce the “prompt” pulse height while simultaneously producing two exponential tails, one due to forward scattering (which lasts a few seconds), as well as a much longer tail due to diffuse scattering.[10] The prompt pulse (“ballistic” photons) is unscattered (therefore unbroadened in time) and reduced in amplitude. Absorption also reduces the prompt pulse height so that the total surviving fraction is $e^{-\tau}$, where $\tau = \frac{4}{5}\frac{R}{R_E}\log_\mathrm{e}10$ is the total optical depth, as mentioned above. Note that the ~20% extinction is modest for the range considered above (1000 ly), but becomes unmanageable for distances substantially greater than R_E.

Thus in this example, even with beam spreading to cover a 10 AU habitable disk at the target solar system, the laser outshines its parent star, in broadband visible light, by a factor of 5000! Moreover, we must not forget that advanced civilizations are supposed to be more *advanced* than we are! Thus “Earth 2000” technology should be a lower bound on the technical sophistication of extraterrestrial civilizations. With a modest extrapolation of another 2–3 orders of magnitude in delivered flux, which can hardly be considered daring given the Moore’s law pace of the optical laser industry, we conclude that a moderately advanced civilization should have no trouble outshining its parent star by six or more orders of magnitude.

3. THE TARGETED SEARCH AT AGASSIZ STATION

Based upon the arguments above, and their elaborations (which evolved during a set of workshops sponsored by the SETI Institute in 1997-9), we designed and built a detector system for pulsed laser beacons. It saw first light on 19 Oct 1998, and has run continuously since.

3.1. How it works

Our search rides piggyback on a radial velocity survey being conducted by two of the authors.[11,12] Roughly half the light that reflects off the entrance slit of the echelle (about one third of the total light) is deflected into our photometer, as shown in Figure 1. This light is re-imaged and passes through a beamsplitter into two hybrid avalanche photodiodes (Hamamatsu R7110U-07), whose outputs feed a pair of multi-level discriminators with levels corresponding to roughly 3, 6, 12, and 24 optical photons. By time stamping level crossings with a LeCroy MTD-135, we obtain an approximate “waveform” of incoming pulses. Coincident pulses seen in each channel trigger the microcontroller to record the waveform profile and arrival time in the two channels. A “hot event” veto filters out a class of large amplitude, bipolarity signals that appear to be produced by breakdown events in the photodetectors. Counters, and various controls and monitors allow us to test the apparatus and monitor its long term fitness. Fiber-coupled LEDs test the detectors and coincidence electronics before every observation. This photometer is really solid – made from 1.25 cm aluminum, milled and bolted together, it weighs about 30 kg. You can see the guts of it in Figure 2.

The diagnostic data, along with coincident pulse data, are sent to a PC and recorded in a log file. After each night of observations, the log files are automatically transferred to computers at Harvard University where they are

---

[*] “Helios” is a diode-pumped Yb:S-FAP laser in development at LLNL for inertial confinement fusion that is potentially capable of 3 ns, 3.7 MJ pulses (10^{15} W) at 349 nm (or 4.7 MJ at its native 1.047 μm wavelength) at ~10 Hz rep-rates.[9]

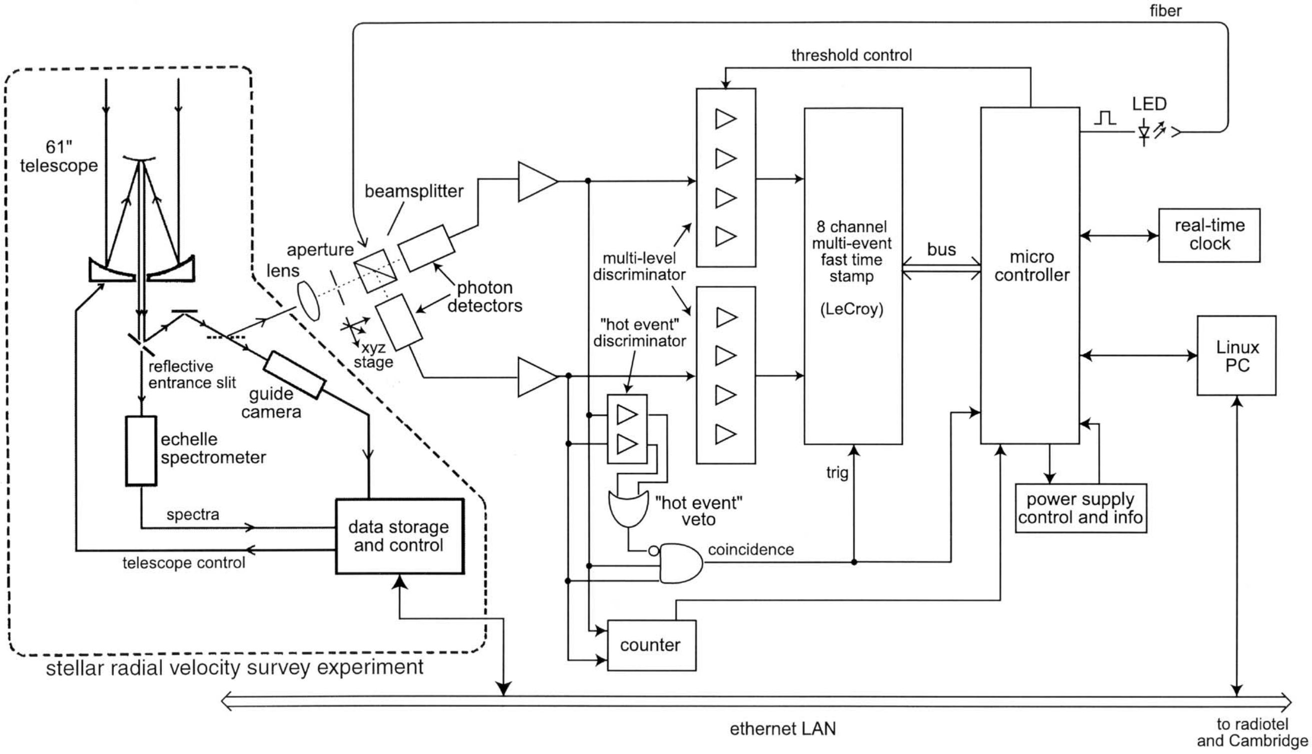

Figure 1. Block diagram of the Agassiz Station targeted optical pulse detector. Unused light from the echelle spectrograph is imaged onto a pair of hybrid avalanche photodetectors, whose coincidence triggers fast time-stamping of waveform crossings through four preset levels.

incorporated into a web-enabled database to facilitate analysis. We track the data through automated daily emails which summarize the previous night's observations. Additionally, the web-enabled database allows us easily to view the data in many forms: chronological summaries, ordered searches by various criteria, observational summaries for individual objects, diagnostic data for particular observations, etc.

Our target list is composed of objects being surveyed both for SETI and for other astrophysical interests. Dave Latham and colleagues have recently begun characterizing 11,000 F, G, and K dwarfs (2000 completed thus far) for possible observations by next generation targeted SETI searches. Specifically, they're looking for evidence of stellar companions that would interfere with planets in the habitable zone. A sample of G dwarfs is being probed for close-in giant planetary companions to determine their galactic frequency and metallicity distribution. Various other programs observe a variety of other targets (A dwarfs, very young stars, very old stars in the Solar neighborhood, etc.).

3.2. Results of two years' operation

From October 1998 through mid-December 2000 (manuscript deadline!), the targeted search has performed nearly 20,000 observations of nearly 5,000 stars, for a total of 2000+ hours of observation.

During this time (excluding the months of high humidity, when the detectors exhibit abnormally high backgrounds; this excludes roughly 30% of the data) we have had ∼700 "hits." We define a hit to be an instance when the lowest thresholds are simultaneously exceeded in both channels. Although all hits are recorded, the "waveforms" are automatically passed through a filter which enforces certain sanity checks: the signals seen in each channel must be roughly the same amplitude (within one level of each other), and they must overlap in time (this is used to exclude a class of hits in which one channel rises again after the other channel shows no signal). The subset of hits which pass this test are labeled "good hits"; to date, we have registered 173 good hits. We do not believe that this categorization

Figure 2. The targeted search photometer, with covers removed. Light enters from the rear of the righthand compartment, focused onto a 15 arcsec aperture, then passes through a beamsplitter onto the pair of HAPDs on their 3-axis stage. The detectors run at a gain of $\sim 4 \times 10^4$, producing $\sim$50 μV negative pulses into 50 Ω, which are amplified and sent to the electronics in the lefthand compartment. The latter perform coincidence, 4-level ADC, timing, logging, hot-event veto, and communication with the host Linux PC. The photometer measures 25×25×60 cm, and weighs 30 kg.

scheme misses extraterrestrial beacons: the LED test flashes, which are done before every observation, have *never* failed this test.

A summary of the database to date reveals that the 173 good hits are distributed over 148 separate targets (with 3099 targets exhibiting no good hits at all), with 132 of them exhibiting a single good hit; the average observation time of each of these (over the entire observational period) is 1.6 hours. The 16 targets registering more than one good hit all had significantly longer total observation times: the 11 targets with 2 good hits averaged 2.7 hours, and the 5 targets with 3–5 good hits averaged 6.7 hours each. The latter were reobserved intentionally, in search of evidence suggesting optical bursts, clusters, or periodic flashes.

Our conclusion to date is this: Among the good hits we have no evidence of clustering or periodicity from any candidate star; the events are distributed impartially among the targets observed, given the low hit rates and corresponding small-number statistics. There is additionally no correlation of hit rate with stellar magnitude, confirming the conclusion that Poisson doubly-coincident "accidentals" do not contribute candidate events at ordinary single-photon count rates. From the results so far, therefore, we conclude that we have found no evidence for pulsed optical beacons from extraterrestrial civilizations.

In considering this conclusion, one must keep in the back of one's mind the possibility that a transmitting civilization might choose to send a solitary pulse, or, equivalently for our observational protocol, a pulse repetition rate less than, say, once per hour. To put it another way, what do you do with 150 isolated non-repeating events... particularly when any one of them, if authentic, would constitute the greatest discovery in the history of humankind?! You find a better way to do the experiment. That is the subject of the next section.

4. SYNCHRONIZED OBSERVATIONS

Given our current background level of roughly one "good hit" per night of observation, a single optical pulse from an extraterrestrial civilization would likely be dismissed as a background hit. To attract attention, the signal would have to be composed of successive pulses from a source candidate, perhaps exhibiting non-random arrival times. As we remarked above, we recognize that this is a shortcoming of the experiment – we may miss a true beacon.

To address this problem, we are collaborating with David Wilkinson and colleagues at Princeton University to duplicate our experiment on their 0.9 m Cassegrain telescope in the Fitz-Randolph Observatory. This telescope will follow the Harvard telescope through its nightly observing programs, synchronized via the internet; simultaneous observations should begin early this year.

We have added an absolute time-tagging daughterboard to both systems: It is reset by the 1 pps pulse from GPS-disciplined clocks at each observatory, it is clocked at 10 MHz, and its 24-bit count is latched upon receipt of a coincidence. Thus we read the absolute time to $\Delta\tau = 100\,\mathrm{ns}$. Given that the baseline between observatories is approximately $L/c = T = 1.6\,\mathrm{ms}$ of light-time, the timing precision not only permits us to identify approximate coincidences; it actually defines an error ellipse in the sky whose narrow dimension is of order $\Delta\theta \approx \Delta\tau/T = 12\,\mathrm{arcsec}$. This is comparable to the observed target field, as set by the 15 arcsec aperture stop. Thus with good accuracy we can verify that a candidate two-observatory coincident event is consistent with the observing geometry.

To see how effective such a scheme is in eliminating uncorrelated events at the two observatories, imagine a "good hit" rate r_b of each experiment of 1 hit per hour (this is a factor of 10 higher than we observe in practice), and let us require that each candidate event pair be within a broad time window of, say, ΔT=1 ms; then the combined background rate is $r_{\mathrm{both}} = r_b^2 \Delta T = 3 \times 10^{-7}$ hits/hour, or 1 hit every 300 years. With such a low background rate, we would have to examine seriously the astrophysical and extraterrestrial significance of even a single coincidence at the two observatories.

5. ALL-SKY OPTICAL SETI SURVEY

These targeted searches have a significant shortcoming – after two years of data collection, we have covered less than one millionth of the sky. With $\sim 10^6$ sun-like stars within 1000 ly, and the possibility that advanced life may exist in the voids between stars, a complementary observing strategy of targeted searches and sky surveys represents the greatest chance for success in optical SETI.

With support from The Planetary Society we have begun construction of a wide-field telescope with a fast, pixelated photodetector whose sensitivity to pulsed optical beacons is complementary to our targeted search. The telescope will view a $1\overset{\circ}{.}6 \times 0\overset{\circ}{.}2$ field with a pair of 512-pixel photodetectors operating in meridian transit mode, with adjustable declination and fixed hour angle: each pixel will cover 2.3 square arcmin. It will scan the sky once in 1200 hours – roughly 150 clear nights – with a minimum viewing time of 48 seconds per target; polar regions are viewed for longer periods of time.

This all-sky survey will look for pulsed optical beacons with a strategy complementary to the targeted approach. In the latter experiment, we are able to choose stars that we believe are likely to harbor life (as well as the targets in the radial-velocity surveys), and observe them for many tens of minutes. The all-sky survey will observe these stars, and millions more, but for shorter periods of time. Freeman Dyson has remarked that the SETI community's bias towards observing stars may even be misplaced; extraterrestrials may live in, and transmit from, the *voids* between the stars. The all-sky survey will observe these areas too. Although low-duty-cycle optical beacons may be missed in the all-sky survey, they are guaranteed to be on the target list (assuming that they're visible from the northern hemisphere). As with SETI at all wavelengths, we believe that a balanced strategy of careful observations of candidate stars coupled with broad surveys of the entire cosmos represents the best chance for contact.

We summarize in the following subsections the architecture and implementation of this challenging "all-sky OSETI."

5.1. Telescope and Physical Infrastructure

The 1.8 m telescope for the all-sky survey is being fabricated by fusing a thin glass slab over a spherical form, then figuring and polishing. Including the effects of spherical aberration, the primary mirror's star images will exhibit

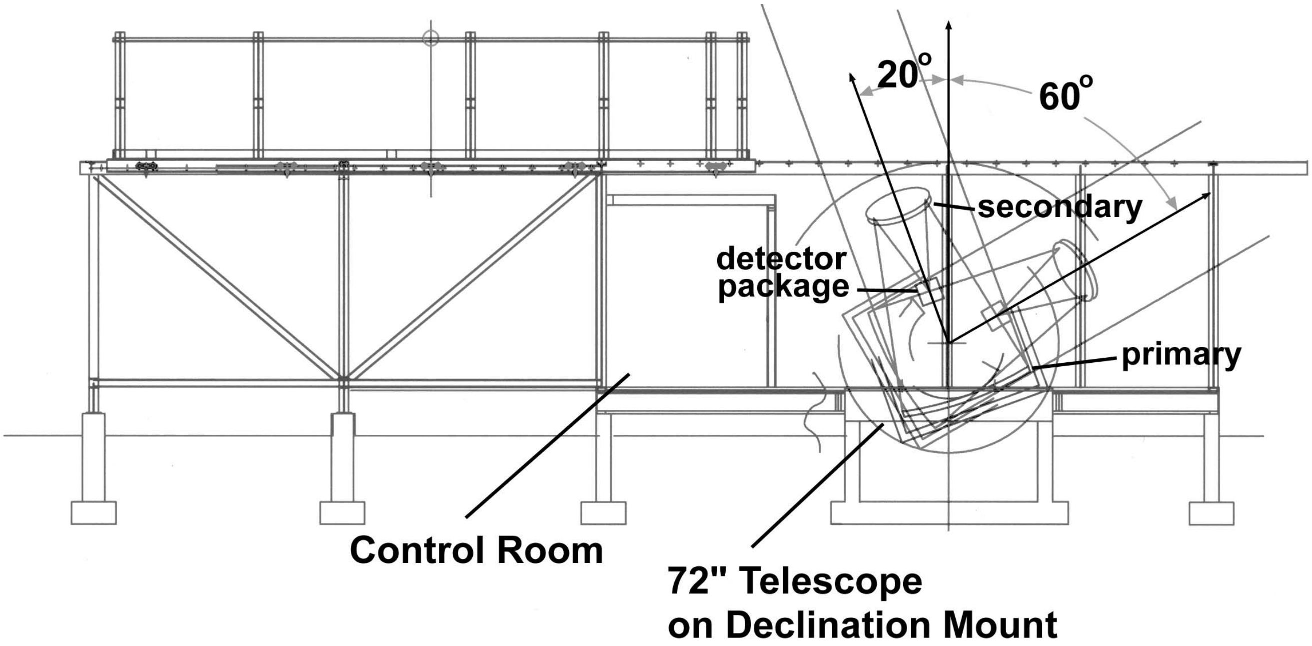

Figure 3. Rolloff-roof "Garage-Mahal" observatory for The Planetary Society's all-sky optical SETI. The meridian-transit survey covers a 1°.6 declination stripe, repositioned daily. The central section of the south wall rolls down, so the 1.8 m pseudo-Newtonian spherical light bucket can view the southern sky.

~ 1.6 arcmin of blur.[†] The telescope (perhaps better described as a "light bucket," given its limited optical quality) is a "pseudo-Newtonian" with a 0.9 m flat secondary mirror at 22.5° with respect to the optical axis. The detector optics consists of a large beam splitter and three folding mirrors to image the beams onto two rows of detectors attached to a single printed circuit board. The detector array consists of two sets of 8 multi-anode photomultiplier tubes (each with 64 anodes), for a total of 1024 pixels; each of the two sky images is observed with an array of 8×64 pixels, linked to an array of coincidence and timing circuitry similar to that of the single-pixel targeted search described previously.

The dedicated observatory structure measures 9 m (N-S) × 5 m (E-W), and is shown in Figure 3. The telescope sits on an isolated concrete pier in the southern part of the building, while the northern part is comprised of an enclosed control room for electronics, computers, motors and other equipment. A rolling roof is suspended on rails that span the length of the building and extend 7 m farther to the North on a steel support structure. This roof is parked to the North during observations; it can also extend slightly over the South wall so that heavy equipment can be lifted into the building using a steel beam in the roof truss structure. A section of the South wall rolls down for viewing southerly declinations.

5.2. Multi-pixel Photomultiplier Tubes

Without the recent advances in the photodetection industry, it would be difficult to pursue this all-sky optical survey. Hamamatsu, among others, recently introduced multi-pixel photomultiplier tubes, which have the radiant sensitivity (quantum efficiency of ∼10-20% for 300-550 nm), gain ($\sim 10^6$), and the speed (rise time of ∼1 ns and FWHM of ∼3 ns) of traditional, single-pixel photomultiplier tubes in each of 64 independent pixels. The tradeoff for adding pixels is some anode nonuniformity and crosstalk. These tubes, like all photomultipliers, also show poor pulse height resolution owing to the statistics of cascaded, low gain stages; individual photoelectrons are incompletely

---

[†] A parabolic primary mirror has severe off-axis coma for the wide field and small f-number that our experiment requires. A fast parabola without additional correctors is worse than a sphere, and harder to make.

discriminated from unresolved pulses of, say, five photoelectrons, but are easily separated from ten photoelectron pulses. The HAPDs used in the targeted search, by contrast, have excellent pulse height resolution.

A pair of pixels – one from each declination stripe – observing the same $1\rlap{.}'5 \times 1\rlap{.}'5$ patch of the sky will function like the single-pixel experiment described earlier. Each pixel pair produces a pair of nanosecond speed analog electrical signals that must be constantly monitored. With 512 such pairs, the electronics to detect large amplitude coincident pulses becomes complex.

5.3. OSETI Signal Processor

Analog data from these 512 pairs of pixels will be analyzed in thirty-two identical full-custom chips – the signals from two detectors are handled by four chips. These chips are presently being designed in-house, and will be fabricated on TSMC's $0.25\,\mu$m process. Their primary purpose is to monitor pixel pairs for large amplitude pulses; additionally, the chips will give a crude image by measuring the single photon count rates for each pixel.

The operation of the signal processor chip goes like this: each PMT anode drives a 500 Msps 7-level flash converter, whose outputs are used to sense a coincidence between corresponding pixel pairs. Such a coincidence triggers steering circuitry that loads a 2K-deep shift register with $4\,\mu$s of waveform data from those two pixels, along with a short length of pre-trigger waveform buffered in a short shift register that precedes the steering MUX. An event flag signals a central microcontroller to download the event's data, via a leisurely bus with conventional interlocked handshaking.

This signal processor chip includes an "astronomy" function, which gates counts exceeding a selectable threshold corresponding to any pair of selected anodes. Internal counters then measure rates during a programmable length of time, and pass this measure of the light intensity on the pixel pair to the outside world. By cycling through the various pixels, the microcontroller can obtain a crude image.

5.4. Other Electronics

All thirty-two signal processing chips, along with the 16 photomultipliers, reside on a single PC board. A microcontroller orchestrates the flow of data between the signal processors and a serial port connection to an external computer. It also monitors pixel countrates and performs various diagnostic functions.

5.5. Software

Fast electronics, replicated into a wide array, produces a firehose of data: With 1024 7-bit flash converters operating at 500 MHz, the pre-filtered data rate is 3.5 *tera*bits per second. That's the contents of all books in print, every second! The signal processing chips reduce the data flow to a manageable rate; they do this by comprising one of the world's largest garbage cans. It is the event of a coincident trigger between corresponding sky pixels that extracts a short burst of the internal data flow to the patiently waiting downstream microcontroller, thence to its larger server PC. Nevertheless, we can expect substantially more data than with the targeted search since we are viewing a much larger portion of the sky.

The software, like the chips, has two primary functions: astronomy, and coincidence detection/archiving. It monitors the pixel countrates to obtain a sky image, which is used to determine accurately where the telescope is pointed. This allows it to tag the locations and times of coincidences. Additional software distills the realtime data to a searchable archival database.

5.6. Hole in the ground

At this time we have the concrete poured, the telescope blank in process, and the signal processor chip nearing completion. We hope to see first light this year; a realistic estimate puts full operation into the following year. Stay tuned!

ACKNOWLEDGMENTS

We gratefully acknowledge the enlightened and continued support of The Planetary Society (TPS) and the Bosack-Kruger Charitable Foundation, and additional support from the SETI Institute. David Brown's matching contribution to TPS made possible the all-sky project. We thank Nathan Hazen and James Oliver for designing the "Garage-Mahal" observatory building. We are indebted to the night observers at Agassiz Station for their tireless efforts that bring us so many photons.

REFERENCES

1. S. Kingsley, ed., *Proceedings of the Search for Extraterrestrial Intelligence (SETI) in the Optical Spectrum*, vol. 1867, SPIE, 1993.
2. S. Kingsley and G. Lemarchard, eds., *Proceedings of the Search for Extraterrestrial Intelligence (SETI) in the Optical Spectrum II*, vol. 2704, SPIE, 1996.
3. R. Schwartz and C. Townes, "Interplanetary and interstellar communication by optical masers," *Nature* **190**, pp. 205–208, 1961.
4. M. Ross, "Search via laser receivers for interstellar communications," *Proc. IEEE* **53**, p. 1780, 1965.
5. C. Townes, "At what wavelengths should we search for signals from extraterrestrial intelligence?," *Proc. Natl. Acad. Sci.* **80**, pp. 1147–1151, 1983.
6. A. Howard et al., "Optical SETI at Harvard-Smithsonian," in *Bioastronomy '99 - A New Era in Bioastronomy*, G. Lemarchand and K. Meech, eds., *ASP Conference Series* **213**, pp. 545–552, 2000.
7. A. Howard, P. Horowitz, and C. Coldwell, "An all-sky optical SEII survey," in *Proceedings of the 51st IAF Congress in Rio de Janeiro, Brazil*, *Acta Astronautica* , in press, 2000.
8. A. Howard and P. Horowitz, "Is there 'RFI' in pulsed optical SETI?," in *The Search for Extraterrestrial Intelligence (SETI) in the Optical Spectrum III*, S. Kingsley, ed., *SPIE* , this volume, 2001.
9. W. Krupke, "Diode-pumped solid state lasers for ife," in *2nd Annual International Conference on Solid State Lasers for Applications to ICF*, Commissariat a l'Energie Atomique (CEA), Paris, France, 1996.
10. J. Cordes, "Pulse broadening in the optical and IR from interstellar grains," in *Proceedings of the Nov. 15 "Optical SETI Workshop", Space Sciences Lab, U.C. Berkeley*, SETI Institute, 1997.
11. D. W. Latham, "Surveys of spectroscopic binaries at the center for astrophysics," in *Complementary Approaches to Double and Multiple Star Research*, H. A. McAlister and W. I. Hartkopf, eds., *ASP Conference Series* **32**, p. 110, 1992.
12. D. W. Latham, "Target selection for planet searches," in *A New Era for Bioastronomy*, G. Lemarchand and K. Meech, eds., *ASP Conference Series* **213**, p. 137, 2000.