Available online at www.sciencedirect.com
ScienceDirect
PERGAMON
Acta Astronautica 61 (2007) 78–87
ACTA ASTRONAUTICA
www.elsevier.com/locate/actaastro

Initial results from Harvard all-sky optical SETI

Andrew Howard[a,*], Paul Horowitz[a], Curtis Mead[a], Pratheev Sreetharan[a], Jason Gallicchio[a], Steve Howard[a], Charles Coldwell[a], Joe Zajac[b], Alan Sliski[a]

[a]*Physics Department, Harvard University, Cambridge, MA 02138, USA*
[b]*Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA*

Available online 16 April 2007

Abstract

We report initial results from the Harvard/Planetary Society all-sky search for pulsed optical signals from other civilizations, which saw "first light" on 11 April 2006 after 6 years of planning and construction. To survey the northern sky ($-20° < \delta < +70°$), our 1.8-m spherical $f/2.5$ optical telescope images a $1.6° \times 0.2°$ patch of sky on two matched focal planes with a total of 1024 photomultiplier tube pixels. Each pair of pixels images the same 2.3 square arcminute patch of sky, and fast electronics filters the incoming visual band light for nanosecond pulses. Coincident optical flashes in a pair of pixels triggers one of 32 PulseNet full-custom chips to record the pulse profiles at nanosecond resolution. This experiment is an all-sky, kilopixel evolution of our targeted search. Its meridian transit survey mode requires $\sim$ 150 clear nights to cover the northern sky with 1-min dwell time per source point. Focus in this talk will be given to describing the capabilities of the all-sky search instrument, the first few months observations, and constraints that these observations place on the density of pulsed optical signals in the galaxy.

Keywords: Extraterrestrial intelligence; Astrobiology; Instrumentation; Detectors

1. Introduction

We report initial results from the Harvard all-sky optical search for extraterrestrial intelligence (SETI). This paper is a companion to "Challenges in the first all-sky optical SETI" by Mead et al. [1], which also appears in this volume. The reader is strongly encouraged to read that paper first since it describes the all-sky telescope, observatory, and instrument, as well as other experimental details. Similar descriptions can also be found in [2]. Readers are refered to [2–4] for general descriptions of SETI and optical SETI, respectively.

This paper focuses on measurements made to calibrate the all-sky telescope and instrument, and on an initial set of observations and their implications. It extends to sky surveys the optical SETI data analysis techniques developed in [4] for targeted observations. The result of this analysis is a constraint on the density of pulsed optical transmitters in the galaxy, as a function of transmitter repeat time. The all-sky observations reported here total 17 h over three nights. While these totals are small compared with the targeted search ($\sim$2400 h over 5 years), the wide field of view and large number of pixels mean that it has already observed significantly more stars and sky area than the Harvard targeted optical SETI (see [4]). The initial observations also provided a way to develop general observing procedures and to plan for automated observations.

* Corresponding author.
E-mail address: andrew@alum.mit.edu (A. Howard).

0094-5765/$ - see front matter © 2007 Published by Elsevier Ltd.
doi:10.1016/j.actaastro.2007.01.038

2. Calibration and sensitivity measurements

2.1. Calibration of telescope position

Before observing, it is important to know precisely where in the sky the telescope is pointed. The all-sky telescope and camera were calibrated by observing a transit of the Moon. In the lunar transit depicted in the two panels of Fig. 1, the center of the Moon crosses the local meridian with a declination of d_0 at a time t_0. (The right ascension $\alpha_0 = t_0$ because it is a meridian transit.) As depicted in the left panel of Fig. 1, the Moon's leading limb grazes the left edge of a photomultiplier tube (PMT) at time t_1. Measurement of t_1 yields the offset between the midline of the PMT array and the local meridian via Eq. (1), which is derived by inspecting Fig. 1:

$$\Delta t = t_1 - t_0 + t_M - t_{PMT}, \quad (1)$$

where t_M (Moon's radius) and t_{PMT} (the distance from the PMT edge to the vertical midline of the PMT array) are measured in drift time.

A similar analysis of the telescope declination d_1 (which is set through the telescope drive controller) that aligns the bottom limb of the Moon with the top of a PMT yields the offset between the horizontal midline of the array and the telescope declination:

$$\Delta d = d_1 - d_0 - d_M - d_{PMT}, \quad (2)$$

where d_M is the Moon's radius.

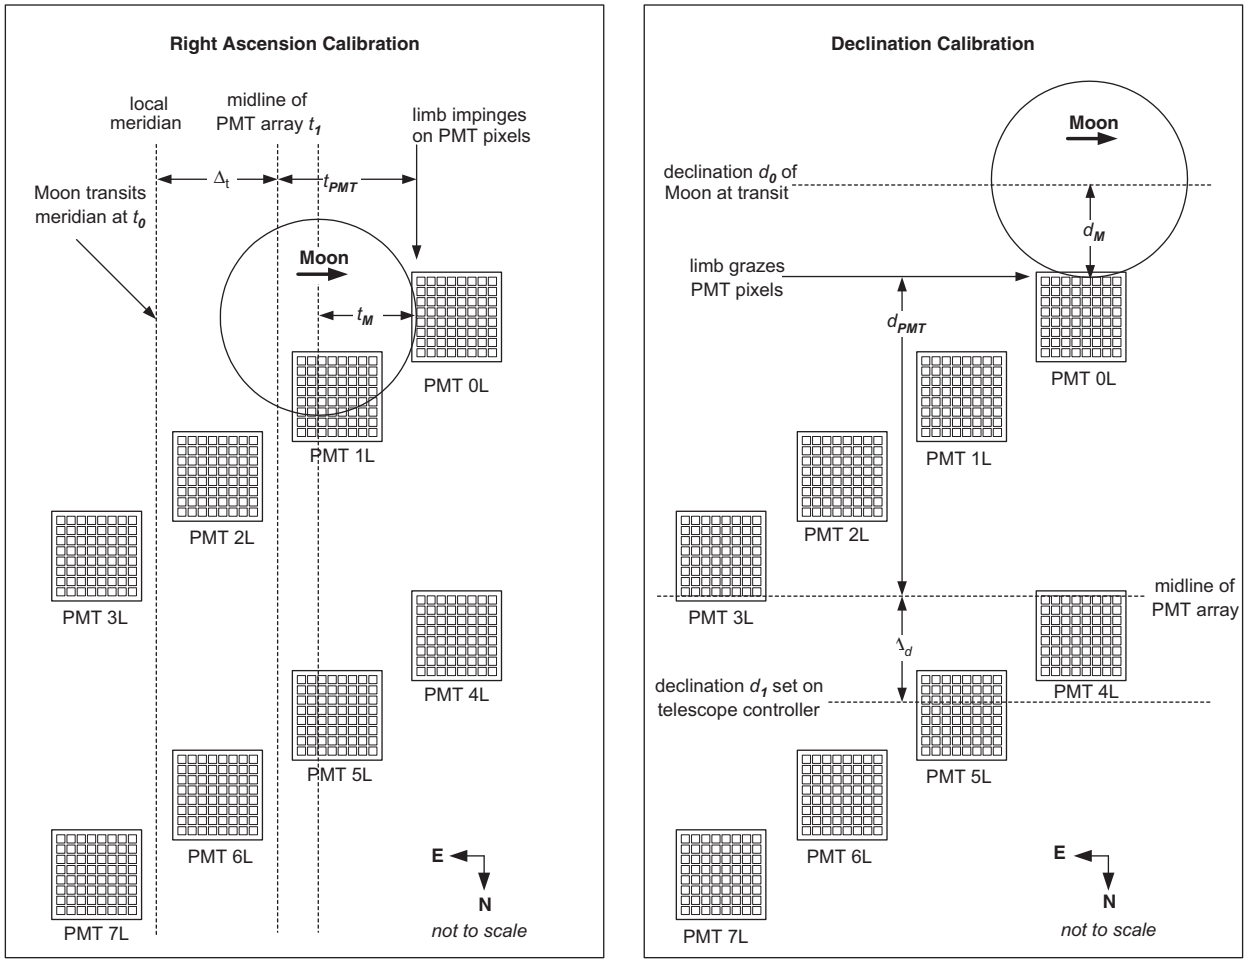

Fig. 1. Calibration of telescope position using PMT 0L. Each panel depicts the Moon moving through the eight PMTs in the left focal plane. The left panel shows calibration of right ascension (α) by measuring the time t_1 that the right edge of the Moon impinges on the edge of the PMT 0L. The right panel shows calibration of declination (δ) by recording the telescope declination d_t for which the bottom edge of the Moon impinges on the top edge of PMT 0L.

These measurements were performed on PMT 4L and PMT 5L for the lunar transit on 7 May 2006 yielding $\Delta t = 3^{\mathrm{m}}13^{\mathrm{s}}$ (=0°48′) and $\Delta d = -1°5'$. In physical dimensions, the vertical midline of the PMT array is 63 mm to the left of the image of the local meridian (for the left PMT array; it is mirrored in the right PMT array), and the horizontal midline of the PMT array is 85 mm above its expected location based on the telescope drive setting.

In the future, calibrating the telescope position may be an automatic function done photometrically using PulseNet's astronomy mode to watch stars drift through the PMT arrays.

2.2. Minimum observing declination

Measuring the southern-most telescope observing angle determines the fraction of the sky accessible to the all-sky instrument. During normal operation, a mercury limit switch prevents the telescope from pointing below $\delta = -20°$. At this angle, the telescope points significantly above the tree line at Oak Ridge Observatory and has an unobstructed view of the sky. By removing the bolted-on south wall of the observatory building (below the barn doors) and rotating the south limit switch to accommodate, there is nothing to prevent the telescope from observing down to $\delta \approx -30°$.

Thus, the telescope is capable of observing declinations in the range $-30° < \delta < +90°$. The area of this portion of the sky is

$$\frac{1}{4\pi}\int_0^{2\pi}\int_{-30°}^{90°} \cos\delta \,\mathrm{d}\delta\,\mathrm{d}\alpha = 75\%. \quad (3)$$

This is a slight increase in sky area from the 64% for the planned declination range $-20° > \delta > +70°$.

2.3. Camera sensitivity

The sensitivity of the all-sky instrument is determined largely by the threshold voltage used to trigger coincident PMT signals, e.g. $V_{\mathrm{thresh}} = \mathtt{V_{bias}} - \mathtt{V_{ref}[1]}$, where $\mathtt{V_{bias}}$ is the input bias voltage on all amplified PMT signals and $\mathtt{V_{ref}[1]}$ is the voltage that PMT signals are compared to. Unfortunately the camera is not noise-free. There are several reasons that the PulseNets may trigger on smaller PMT signals (because of threshold/bias offsets and asymmetries) or other signals may couple into the PMT signals (before or after amplification). The list of candidates for such erroneous triggers include:

- Coupling with clock traces on the daughterboards.
- Coupling with corona discharge on the PMTs.
- Offset-voltage differences between PulseNet samplers—the $1-\sigma$ variation for one sampler is 11 mV.

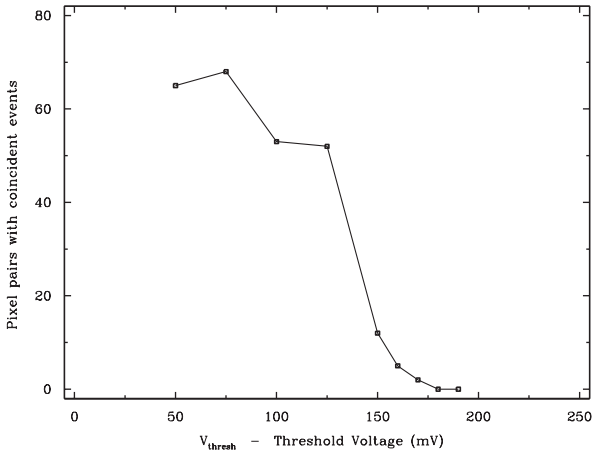

Fig. 2. Number of pixel pairs (out the 512 total pairs in the PMT pixel array) with coincident events during 1-min observations as a function of threshold voltage. Compare with Fig. 3, which plots these measurements sorted by PulseNet pixel number.

- Amplifier gain variations—the specification for small signal gain at 900 MHz[1] is 21^{+3}_{-2} dB. This is probably a conservative estimate, but on the high side of the range, 3 dB is a ∼ 40% increase in gain.
- Variable PMT cathode sensitivity—the variation can be as high as a factor of 2–3 on a tube, but typical pixels vary by ∼ 30–50%; note that this variation will not affect the single photoelectron pulse size, just the rate of such pulses.

To determine the threshold voltage to be used in observations and to check for sources of noise, the camera sensitivity was measured. Each measurement was a 1-min observation on a dark night during which the number of pixel pairs (for all 512 pixel pairs in the array) that registered at least one coincident event was counted. Many observations were made at several threshold voltages to obtain the plot in Fig. 2.

Note that above $V_{\mathrm{thresh}} = 170$ mV, no coincident events were recorded. Between 125 and 150 mV, Fig. 2 shows a sharp increase that levels off to 50–60 pixel pairs for $V_{\mathrm{thresh}} < 125$ mV. For $V_{\mathrm{thresh}} < 150$ mV, the total number recorded in each 1-min observation was approximately the same. This is likely the result of a finite event processing rate (a few per second) with the PC104 and host computer. For observations at the high end of the threshold range ($V_{\mathrm{thresh}} \geqslant 150$ mV), the events were concentrated in one to three pixel pairs in the entire array, each of which received 20–50 events.

[1] Typical gain at lower frequencies is about 1 dB (12%) lower, but minimum and maximum values are not stated in the data sheet.

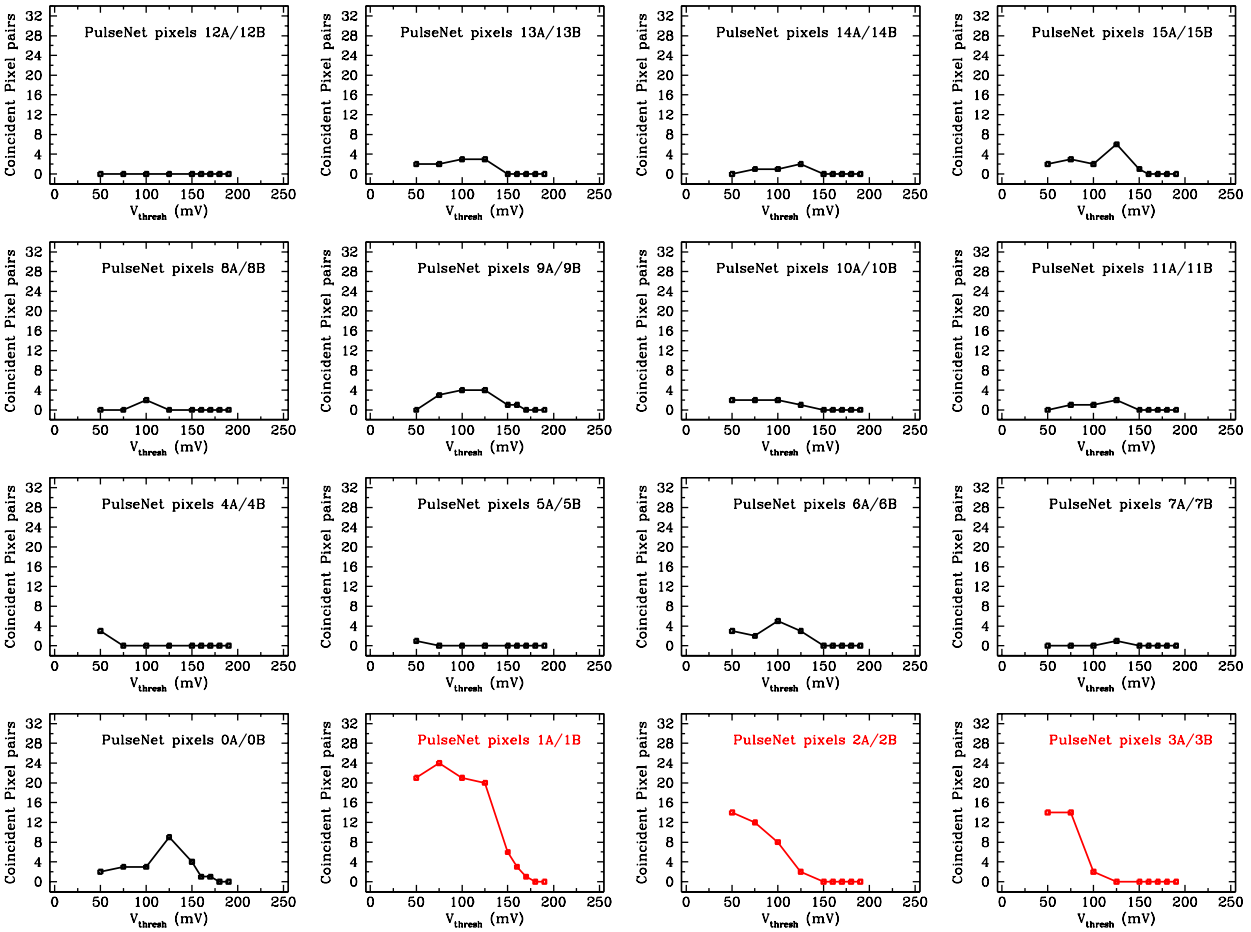

Fig. 3. Dependence of PulseNet pixel number on sensitivity. For each numbered PulseNet input pixel pair (of which there are 16 per PulseNet numbered 0–15), a plot shows the number of pixel pairs (from all 32 daughterboards) with coincident events during 1-min observations as a function of threshold voltage. The three red traces (the three right-most plots in the bottom row) highlight an excess of coincident pixel pairs in PulseNet pixels 1–3. Compare with Fig. 2 which combines these 16 plots.

For observations with $V_{\text{thresh}} < 150$ mV, the events were distributed much more evenly among the larger number of pixel pairs that recorded any events; in this threshold range, there were no pixel pairs that recorded more than five events.

Fig. 3 plots the same data sorted into 16 subplots according to the PulseNet pixel number.[2] Note the traces for PulseNet pixel numbers 1–3 (highlighted in red); for $V_{\text{thresh}} < 150$ mV, these pixels are responsible for a large fraction of the total events. There is clearly something unusual about these PulseNet pixel numbers 1/2/3, whose corresponding PMT pixels uniformly tessellate the entire array of PMTs. This correlation only appears when the pixels are sorted by PulseNet pixel number, and not by PMT pixel number.

The above data suggest that there are (at least) two classes of the events in Figs. 2 and 3. For $V_{\text{thresh}} \leqslant 150$ mV, the dominant source of events must be spatially localized on the daughterboards in such a way that favors certain PulseNet pixels. The likely culprit is coupling from `fastclock`, the signal that sets the PulseNet sampling rate. Inspecting the daughterboard PCB reveals that signal traces for PulseNet inputs `in00B`, `in01B`, `in02B`, and `in03B` run close to termination resistors for `fastclock`; the trace for

[2] The 512 pixel pairs in the array can be classified by the PMT number (where the photons are detected) or by PulseNet number (where signals from the PMTs are terminated). PulseNet pixel numbers correspond to input pairs (e.g. `in00A/in00B`) on PulseNet chips. PMT pixel pairs with the same PulseNet pixel number may be from quite different locations in the photodetector array, but use the same PulseNet inputs (on different chips) and sample similar electrical environments on their respective daughterboards as the signals are routed from the amplifiers to the PulseNets.

in01B (the pixel with the largest number of pixels with coincident events in Fig. 3) runs *between* the fastclock termination resistors.

There is another hypothesis for the concentration of events in low PulseNet pixel numbers. It is more complicated and does not fit the observed data nearly as well as the fastclock-coupling hypothesis. This hypothesis has to do with the priority given to multiple coincident events that occur on the same clock edge in a single PulseNet chip. A set of such coincidences could be generated by a transient that affects a large number of inputs. When multiple events are received, the memcontroller module in PulseNet gives priority to sampen cbank0, which corresponds to PulseNet pixels 0–3, over sampencbank1...3. The memcontroller module within membank checks for coincident events from the four sampencbank blocks (by checking $\texttt{coinc\_add}\texttt{a}^{\texttt{a}}_{\texttt{b}}\texttt{0...3[3:0]}$ in their number order 0–3 for a diagram of top-level PulseNet blocks), by latching the address (e.g. coinc_add_a0[3:0]) of the pixel pair with an event. Under normal operation, where events are infrequent and uncorrelated, only 1 bit of this 4-bit address will be high (corresponding to the pixel within that block with the event), and the 1-of-4 address will mask the appropriate bits of the streaming input samples ($\texttt{G}^{\texttt{aA}}_{\texttt{bB}}\texttt{0...3[2:0]}$) so that samples from the coincident pixel pair (e.g. $\texttt{G}^{\texttt{aA}}_{\texttt{bB}}\texttt{0[2:0]}$) are passed out of sampencbank and stored in memory. However, if 2 or more bits of coinc_add_a0[3:0] are high when it is latched, the mask will bitwise OR the samples for the pixels with events, resulting in meaningless data stored in memory.

This hypothesis has two observable consequences. One is that there should be more events in PulseNet pixel number 3 than in 2, more in 2 than in 1, etc.[3] Such a trend is not observed in Fig. 3. The other observable consequence is that waveform data from such events will be meaningless combinations of several waveforms, and, when plotted, will appear atypical. A small number of waveforms do have an atypical appearance (among other things, they appear to lack samples at voltage levels that would have triggered PulseNet); however, there are other possible reasons for this.

In either case, a hypothesis that explains events at low pixel number does not account for all events in Figs. 2 and 3. For $V_{\text{thresh}} \geqslant 150\,\text{mV}$, the events are concentrated in a small number of PMT pixels that are independent of PulseNet pixel number and are consistent from one observation to the next. These hot pixels are likely to be on the tail of the V_{thresh} distribution. This distribution probably has contributions from most of the bullet points above: hot pixels may experience some coupling from fastclock, they are probably on the high side of the PMT cathode sensitivity and amplifier gain distributions, and their PulseNet sampler offset voltages are probably a few σ from the mean.

Improving the sensitivity remains a challenge. The clock-coupling mechanism will likely be reduced in a planned daughterboard replacement since the new daughterboard has flying clocks. The hot pixels may be dealt with by individually tuning the components along those signal paths. Such (time consuming) solutions include verifying solder joints connectivity for components near those signals (particularly on amplifier bypass capacitors), replacing the amplifiers on hot pixels with slightly lower gain versions, and changing the 50 Ω termination resistors to reduce PMT signal amplitude (with the attendant risks of signal reflections).

3. Observing procedures

For the sake of consistency, safety, and meaningful results, it was necessary to develop a set of manual observing procedures (Fig. 4) before conducting the initial observations described in Section 4. An important aspect of this is the tests on the camera functionality done each night, and other set of tests before each 30-min observation. Additionally, telemetry data provide constant feedback on the programming state of the PulseNets, clocks, and programmable voltages, the state of power supplies, and environmental conditions. In an experiment looking for rare events, it is easy to confuse an unresponsive or malfunctioning instrument with one that is making observations with null results.

The "self-coincidence test" verifies that all 32 PulseNets can be programmed and properly detect and record coincident events. This test relies on the trick of reversing the standard ordering of $\texttt{V}_{\texttt{ref}}\texttt{[6:0]}$ and $\texttt{V}_{\texttt{bias}}$ by setting $\texttt{V}_{\texttt{ref}}\texttt{[1]} > \texttt{V}_{\texttt{bias}}$ and setting the PulseNets to trigger on $\texttt{V}_{\texttt{ref}}\texttt{[1]}$. Since all 32 PulseNets will immediately register coincidences, this checks the health of these chips as well as the layers of communication between the user interface and the electronics.

[3] This is because software in the host computer calculates int($\log_2$(coinc_add_a0[3 : 0])) to determine the coincident pixel address. When coinc_add_a0[3:0] is a 1-of-4 address, this function properly identifies the pixel address (0–3) within the given block. However, if coinc_add_a0[3:0] has 2 or more bits high, the result will be the highest pixel number (0–3) that registered an event; information about lower pixel numbers are lost when $\log_2$(coinc_add_a0[3 : 0]) is integer truncated. The software can and should be rewritten to identify and handle such multiple coincidence events.

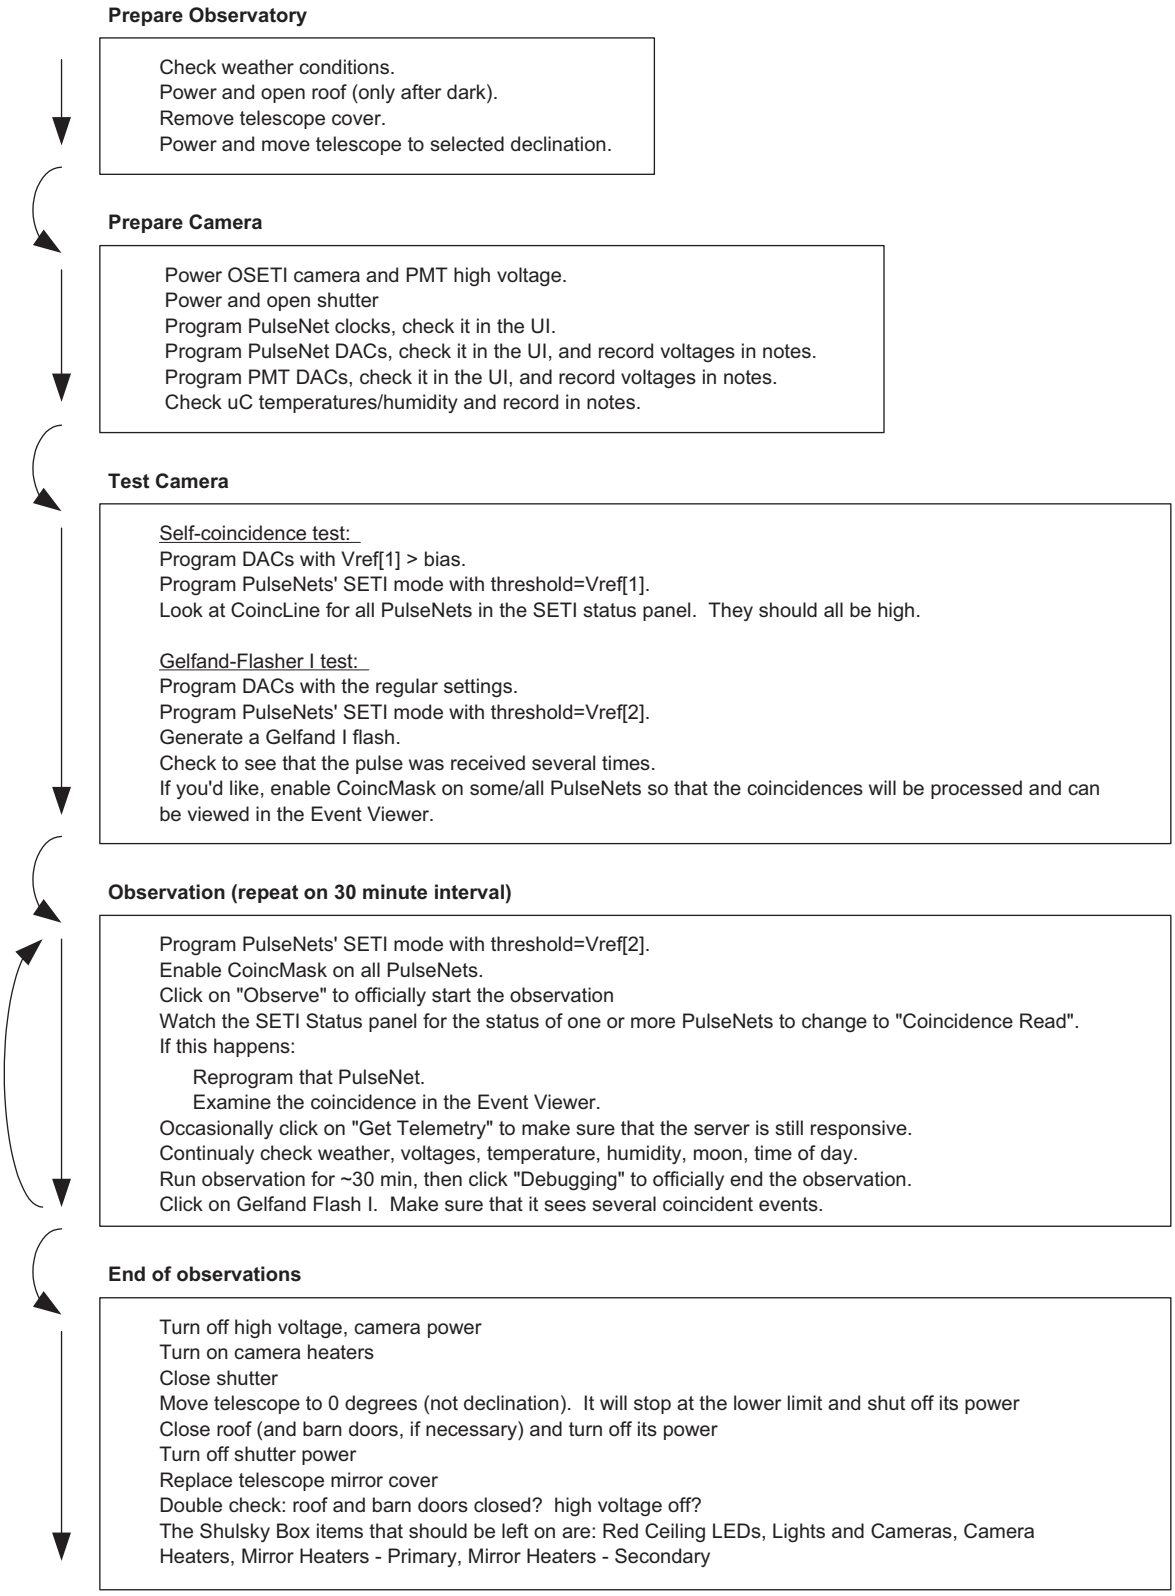

Fig. 4. All-sky observing procedures for initial manual observations.

The Gelfand Flasher I test sends an optical flash at PMT arrays from a defocused LED with pulsed input. It triggers 8–12 simultaneous coincidence events that have a characteristic pulse shape and width (~ 80 ns). It also tests the GPS timestamping electronics; since the coincident events happen simultaneously, their GPS timestamps should be identical.

In order to survey the whole northern sky, these procedures are currently in the process of being automated. We expect to begin automated observations in the fall of 2006.

4. Initial observations

The initial observations listed in Table 1 were conducted over three nights in late April 2006 using the procedures outlined in Section 3. The 34 observations of $\sim$ 30-min duration each totaled 17 h observing time. The observations were all near the equatorial plane so as to maximize sky coverage (at the expense of dwell time per source point).

These observations served several purposes: they tested the instrument, encouraged the development of observing procedures, and, because of the sky coverage and sensitivity of the all-sky instrument, these small number of initial observations comprise the most extensive search for pulsed optical extraterrestrial transmitters ever conducted.

All three nights had completely clear skies, with excellent seeing, and moderate humidity (45–60%). Each night the all-sky instrument was programmed with the following parameters:

- `fastclock` was set to 300 MHz for a double-edge sample rate of 600 MS/s and a sampling interval of 1.7 ns.
- Threshold voltage programming levels: the threshold voltage was 250 mV using $\mathtt{V_{ref}}[2]$ as the trigger with the reference voltages spaced as follows: $\mathtt{V_{ref}}[0] = 1.700\,\mathrm{V}$, $\mathtt{V_{bias}} = 1.500\,\mathrm{V}$, $\mathtt{V_{ref}}[1] = 1.375\,\mathrm{V}$, $\mathtt{V_{ref}}[2] = 1.250\,\mathrm{V}$, $\mathtt{V_{ref}}[3] = 1.100\,\mathrm{V}$, $\mathtt{V_{ref}}[4] = 0.950\,\mathrm{V}$, $\mathtt{V_{ref}}[5] = 0.75\,\mathrm{V}$, $\mathtt{V_{ref}}[6] = 0.500\,\mathrm{V}$.
- PMT high voltages: all PMTs were programmed to 900 V.

There were two "triggers" and zero "events" detected during the first night. No triggers or events were recorded in the last two nights' observations.[4]

[4] Note that the instrument successfully detected the optical pulses in the Gelfand Flasher I tests conducted before each observation. The lack of triggers/events during the observations these two nights does not indicate malfunction.

As with the targeted search, a "trigger" is defined as an instance when one or more PulseNets register coincidence events. The subset of these that pass sanity checks are called "events". The two triggers on 26 April 2006 are not classified as events because in each case the trigger was recorded simultaneously in several PulseNets that correspond to non-adjacent PMT pixels. This non-locality is a signature of background triggers, and is an antisignature of genuine optical pulses imaged through the telescope that are focused to a spot size encompassing at most four adjacent pixels that share a corner. Note that these triggers are probably not due to Čerenkov radiation from cosmic rays. Scaling from [5], the minimum energy cosmic ray that would deliver $\sim 100\,\mathrm{photons/m^2}$ into a PMT pixel is $\sim 10^{16}\,\mathrm{eV}$, which has a frequency about once every several years from an arbitrary point on the ground.

These initial observations covered 1% of the sky (400 sq deg), as illustrated in Fig. 5. Note that this is approximately a factor of 4000 larger than the sky coverage for the entire targeted search. The number of objects observed is also impressive. Using a model of $\sim 10^7$ Sun-like stars in range (and 10^8 total stars in range), 1% sky coverage implies that $\sim 10^5$ Sun-like stars—a factor of ~ 200 more than that in the targeted search—were surveyed for 48 s each. Even with a conservative threshold voltage setting of 250 mV, the initial observations had approximately the same sensitivity as the targeted search ($95\,\mathrm{photons/m^2}$ in $\leqslant 3$ ns), and *no* observations produced events that were consistent with genuine interstellar optical flashes.

5. Implications

5.1. Methodology

By adapting the procedure used in the targeted search (see [4]), we can put constraints on the density of transmitting civilizations in our region of the galaxy. As with that search, let us assume that none of the all-sky observations detected signals from extraterrestrial civilizations (Table 2).

One key difference between the targeted and all-sky data has to do with the definition of an observation. In the targeted search, observations were of individual stars for fixed durations, while all-sky observations are of patches of the sky that drift through the field of view at a certain rate. The all-sky search does not record the names of the stars that it observes, nor would it even be practical to try to enumerate them. The observa- tional target of this search is the sky itself. As such, the methodology for constraining the number

Table 1
All-sky survey—initial observations

Date	Declination range	Observations	Total observing time	Triggers	Events
26 April 2006	0°17′–1°53′	5	2 h 33 min	2	0
28 April 2006	1°53′–3°29′	12	6 h 9 min	0	0
29 April 2006	3°29′–5°05′	17	8 h 19 min	0	0

Note: Initial observations using the all-sky instrument and following the procedures in Fig. 4.

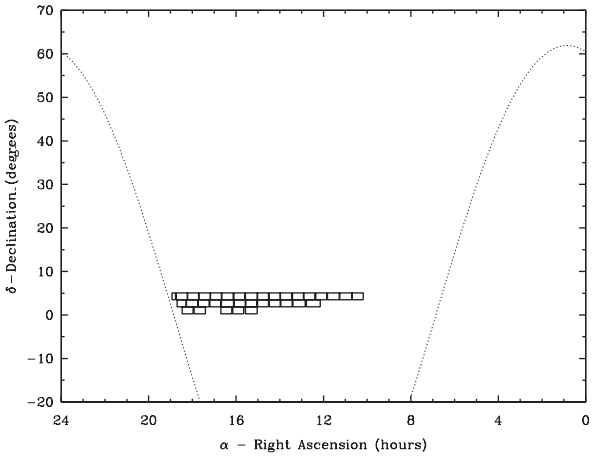

Fig. 5. Initial sky coverage for the all-sky search in equatorial coordinates. The 34 observations over 17 h, each represented by a rectangle in right ascension and declination, are summarized in Table 1. The dashed line represents the galactic equator.

of transmitting civilizations should reflect this distinction between discrete objects and continuous patches of space with densities of objects. In short, the sums must be turned into integrals and the probabilities into probability densities.

A necessary first step in this analysis is constructing models of the density of stars that are close enough for optical communication and of extinction, both of which are functions of the structure of the galaxy and depend on observing direction. In general, the density of Sun-like stars within range is a function of the celestial coordinates, right ascension (α) and declination (δ). Constructing such a detailed model for $\rho_{\text{stars}}(\alpha, \delta)$ is complicated and will be left to future work.

We will use a simple model: a uniform density of stars within range, of which 10% are Sun-like.[5] When integrated over the whole sky, yields 10^8 stars:

$$\int\int \rho_{\text{stars}} \cos\delta \, \mathrm{d}\delta \, \mathrm{d}\alpha = 10^8 \text{ stars}. \tag{4}$$

This model corresponds approximately to a range $R = 1$ kpc and a loss of up to 50–80% of transmitted photons due to extinction in the visible part of the optical spectrum. In this simple model there are $\sim 10^7$ Sun-like stars in range. Note that since there are about 41,000 sq deg in the sky, this works out to $\rho_{\text{stars}} \approx 2400$ stars per sq deg (~ 240 Sun-like stars per sq deg). It is also worth noting that with an instantaneous solid angle of 0.3 sq deg, the all-sky instrument surveys ~ 700 stars (~ 70 of which are Sun-like) at a time.

Although this approximation considerably simplifies the analysis, we will develop a methodology that allows for the use of models of $\rho_{\text{stars}}(\alpha, \delta)$ that depend on celestial coordinates. This will be useful in the near future when such models are developed to analyze data from a full survey.

The next step in the methodology is to represent *each* observation as a function of the celestial coordinates:

$$T_i(\alpha, \delta) = \begin{cases} t_{\text{drift}}(\delta) & \text{for all } (\alpha, \delta) \text{ in the observation,} \\ 0 & \text{otherwise,} \end{cases} \tag{5}$$

where $t_{\text{drift}}(\delta) = (48\,\text{s})/\cos\delta$. That is, when plotted, each $T_i(\alpha, \delta)$ looks like one of the rectangles in Fig. 5, and has an approximately constant value of the drift time through the array for the patch of the sky observed. There are 34 T_i's in the initial observations; to fully cover the northern sky ($-20° > \delta > +70°$) once with $\sim$ 30-min observations, there will be ~ 2700 of them.

Recall that in the targeted search analysis, each observation of duration t_{obsv} has a corresponding probability of detection, $p_{\text{obsv}}(P) = \min(1; t_{\text{obsv}}/P)$, that is, a function of the transmitter repetition period P. The analogous statement in the all-sky search is that for the ith observation, $T_i(\alpha, \delta)$ has an associated probability

[5] For the purpose of this model, "Sun-like" is an F, G, or K dwarf. This may be conservative, particularly if M stars are found to have habitable planets.

Table 2
Comparison of all-sky survey and targeted search data

	Targeted search		All-sky survey	
	DS_{clean}	$DS_{overlap}$	Initial observations	Full survey
Sensitivity (photons/m^2)	100 in $\leqslant$5 ns	100 in $\leqslant$5 ns	95 in $\leqslant$3 ns	17 in $\leqslant$5 ns
Sky coverage (%)	0.0002	0.00005	1	64–75
Objects	4730	1142	$\sim 10^6$	$\sim 10^8$
Observations	11,600	1721	34	$\sim$ 2700
Observations per object	2.5	1.5	1	1
Observing time (h)	1721	244	17	1350–1800
Observing time per object	22 min	13 min	48 s	48–140 s
Events	274	130	0	–
Event rate (h^{-1})	0.16	0.53	0.00	–
Triggers	1066	614	2	–
Trigger rate (h^{-1})	0.62	2.52	0.11	–

Note: Statistics comparing the all-sky survey and the targeted search. Data for the targeted search are from [4]. The events and triggers listed for $DS_{overlap}$ are the number recorded in the Harvard targeted search instrument; none were confirmed by the Princeton instrument. For the all-sky survey, projected numbers are listed for searches ranging from $-20° < \delta < +70°$ to $-30° < \delta < +90°$. The sensitivities listed are for the wavelength ranges for the particular experiments (see [2,4]) and, in the case of the full survey, assume a successful daughterboard upgrade.

density function

$$p_i(\alpha, \delta, P) = \min(1, T_i(\alpha, \delta)/P), \quad (6)$$

where the minimum function bounds $p_i \leqslant 1$.

Although no point in the sky was observed more than once during the initial observations, they will be during longer runs. To fold together these probabilities, we use

$$p(\alpha, \delta, P) = 1 - \prod_i (1 - p_i(\alpha, \delta, P)). \quad (7)$$

The final step in this analysis is to calculate the expected number of signal detections for all observations, S. Recall that in the targeted search analysis, this is just $S = f(P)\sum_{obj} p_{obj}$ and we adjust f so that the Poisson probability of observing zero extraterrestrial signals $e^{-S} = 0.5$. To say it another way, we choose f so that we have a 50% chance of success; we assume that the search was on the cusp of detection and f becomes an upper bound.

With probability densities instead of probabilities, we have to integrate over the celestial coordinates instead of summing over objects:

$$S = f(P) \int\int p(\alpha, \delta, P)\rho_{stars}(\alpha, \delta) \cos\delta \, d\delta \, d\alpha. \quad (8)$$

Using $e^{-S} = 0.5$, as above, and solving for $f(P)$, we obtain an upper bound on the fraction of transmitting civilizations:

$$f(P) = \min\left(1, \frac{\ln 2}{\int\int p(\alpha, \delta, p)\rho_{stars}(\alpha, \delta) \cos\delta \, d\delta \, d\alpha}\right), \quad (9)$$

where the minimum function bounds the density of extraterrestrial transmitters at 100% of stars.

5.2. Constraints on transmitting extraterrestrial civilizations

Applying the methodology of Section 5.1 to the initial observations of 1% of the sky yields the plot of $f(P)$ in Fig. 6 (red trace). For comparison, results from the targeted search (see [4]) are plotted in blue (Harvard search) and cyan (Harvard–Princeton joint observations). The dashed lines from these plots represent the data reported in [4], while the solid lines represent the addition of $\sim$ 1.5 year of observations (through May 2005, when the searches ended). Additionally, the green trace approximates the limits that could be set by null observations with the all-sky instrument by observing the northern sky twice.

The features of the targeted and all-sky searches are apparent in plots of $f(P)$. The wide field of view, 512 parallel pixels, and resulting large number of stars simultaneously observed give the all-sky search the capability of setting more sensitive limits on the fraction of stars with transmitting civilizations. One tradeoff is reduced observation time per star. The $f(P)$ curves for the targeted searches thus extend to the right (longer transmitter repetition periods) at their maximum sensitivities. To be fair, Fig. 6 may underestimate the effectiveness of the targeted searches since they observe a non-random set of stars.

It is important to consider the vertical scale in Fig. 6 in the context of estimates for the number of

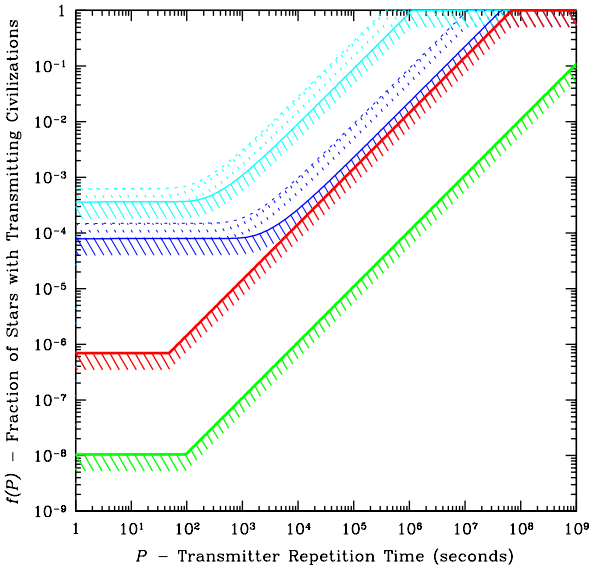

Fig. 6. Upper bounds on the fraction of stars with transmitting civilizations as a function of transmitter repetition time for three searches—the Harvard–Princeton targeted search ($DS_{overlap}$) in cyan (top two traces), the Harvard targeted search (DS_{clean}) in blue (middle two traces), and initial observations from the all-sky survey (1% of the sky) in red (second to bottom trace)—along with projected results from the all-sky instrument after surveying the northern sky twice in green (bottom trace). Two curves are shown for each targeted search: the dashed lines are those reported in [4]; solid lines incorporate data from ~ 1.5 year of additional observations. Each curve asymptotes to $f = 1/N_{obj}$ for small P and cuts off at $f = 1$ for $P \geqslant T/\ln 2$.

transmitting civilizations N, which range from 1 to 10^6. At the optimistic end of the range, $N = 10^6$ corresponds to $f = 2.5 \times 10^{-6}$. If civilizations are packed in the galaxy at that density, the targeted searches are still 2–3 orders of magnitude away from detecting transmitters with relatively short periods (not accounting for the unknown advantage of target selection). The all-sky survey does much better by this metric. Initial observations with the all-sky instrument (1% of the sky) achieved $f(P) \leqslant 7 \times 10^{-7}$ for $P < 48$ s, which corresponds to a limit of $N \leqslant 3 \times 10^5$, using the above model for extinction and stellar density. A full sky survey has the potential to constrain $f(P) \leqslant 10^{-8}$, achieving a sensitivity of $N \leqslant 4 \times 10^4$.

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 and the Shulsky Foundation. William Dally and Gu-Yeon Wei have been invaluable advisors on the design of PulseNet. We recognize the generous grant from the MOSIS corporation for the fabrication of PulseNet. We thank Nathan Hazen and James Oliver for designing the all-sky observatory building and Ray Demarais for his construction of the telescope. Chris Laumann, Chris Clearfield, and David Sliski have also contributed importantly to this project. Portions of this paper were previously published in the first author's Ph.D. thesis [2].

References

[1] C. Mead, P. Horowitz, A.W. Howard, P. Sreetharan, J. Gallicchio, S. Howard, C. Coldwell, J. Zajac, A. Sliski, Challenges in the first all-sky optical SETI, 2006, this volume.

[2] A. Howard, Astronomical searches for nanosecond optical pulses, Ph.D. Thesis, Harvard University, 2006.

[3] J. Tarter, The search for extraterrestrial intelligence (SETI), Annual Reviews of Astronomy and Astrophysics 39 (2001) 511–548.

[4] A.W. Howard, P. Horowitz, D.T. Wilkinson, C.M. Coldwell, E.J. Groth, N. Jarosik, D.W. Latham, R.P. Stefanik, A.J. Willman, J. Wolff, J.M. Zajac, Search for nanosecond optical pulses from nearby solar-type stars, Astrophysical Journal 613 (2004) 1270–1284.

[5] A.W. Howard, P. Horowitz, Is there RFI in pulsed optical SETI?, in: S.A. Kingsley, R. Bhathal (Eds.), Proceedings of the SPIE vol. 4273, The Search for Extraterrestrial Intelligence (SETI) in the Optical Spectrum III, August 2001, pp. 153–160.